ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY

TELECOM ParisTech

2015-ENST-0021

EDITE - ED 130

Doctorat ParisTech

T H È S E

pour obtenir le grade de docteur délivré par

TELECOM ParisTech

Spécialité « Signal et Image »

présentée et soutenue publiquement par

Giovanni CHIERCHIA

le 20 avril 2015

Éclatement épigraphique de contraintes convexes. Application à la restauration d'images, la classification supervisée, et la détection d'images falsifiées.

Directrice de thèse : **Béatrice PESQUET-POPESCU**
Co-encadrement de la thèse : **Nelly PUSTELNIK**

Jury

M. Giovanni POGGI, Professeur, Università Federico II di Napoli	Président
M. Jean-François AUJOL, Professeur, Université de Bordeaux	Rapporteur
M. José BIOUCAS-DIAS, Professeur, Universidade de Lisboa	Rapporteur
M. Alain RAKOTOMAMONJY, Professeur, Université de Rouen	Examinateur

TELECOM ParisTech
école de l'Institut Mines-Télécom - membre de ParisTech
46 rue Barrault 75013 Paris - (+33) 1 45 81 77 77 - www.telecom-paristech.fr

T H È S E

Acknowledgments

I would like to acknowledge my supervisors Beatrice Pesquet-Popescu and Nelly Pustelnik for their support along my PhD thesis and the invaluable guidance on my research topic. I would also like to thank Jean-Christophe Pesquet for his generous support. It was a real pleasure to work with them, to share passion for optimization, to be enrolled into scientific discussions, and to understand how to proceed from an idea to an article. I have learned a lot from them.

My deepest gratitude is due to the members of my jury, Jean-François Aujol, José Bioucas-Dias, Giovanni Poggi, and Alain Rakotomamonjy, for their time. It was a great honor for me to present my work to them.

I would like to acknowledge Telecom ParisTech and Université Paris Est for the great atmosphere in their laboratories. I was lucky to meet there many friends and great young scientists, such as Abdel-Bassir, Andrei, Ania, Audrey, Camille, Caroline, Claudio, Elie, Feriel, Francesca, Francesco, Giovanni, Giuseppe, Hamlet, Irina, Mai, Maxime, Marco, Marwa, Mireille, Mounir, Nello, Paul, Rafael, Raffaele, and Yosra. I would like to thank them all for their kindness, laughters, empathy and advise. After all, the PhD is a human experience too.

Special thanks go to Marco Cagnazzo, Giovanni Poggi, and Luisa Verdoliva, who made it possible for me to join Telecom ParisTech in the first place, but also to pursue my PhD in the best possible conditions.

Finally, I would like to thank my lovely Mireille and my dear family: without them, I would not have been strong enough to complete this thesis, go through all the difficulties, and overcome them.

Abstract

In this thesis, we present a convex optimization approach to address three problems arising in multicomponent image recovery, supervised classification, and image forgery detection. The common thread among these problems is the presence of nonlinear convex constraints difficult to handle with state-of-the-art methods. Therefore, we present a novel splitting technique to simplify the management of such constraints. Relying on this approach, we also propose some contributions that are tailored to the aforementioned applications.

The first part of the thesis presents the epigraphical splitting of nonlinear convex constraints. The principle is to decompose the sublevel set of a block-separable function into a collection of epigraphs. So doing, we reduce the complexity of optimization algorithms when the above constraint involves the sum of absolute values, distance functions to a convex set, Euclidean norms, infinity norms, or max functions. We demonstrate through numerical simulations that the proposed method can efficiently handle constraints based on functions commonly used in image restoration or supervised learning, such as nonlocal total variation, Kullback-Leibler divergence, and logistic regression.

The second part of the thesis presents three contributions grounded on the epigraphical splitting. The first one is a novel regularization for multicomponent images that extends the nonlocal total variation by taking advantage of the structure tensor. The second one is a learning algorithm for efficiently and exactly training a multiclass support vector machine with sparse regularization. The third one is a variational approach to detect image forgeries by using the photo response non-uniformity (a deterministic pattern noise that uniquely identifies each individual camera). We carried out numerical experiments for each application in order to illustrate the efficiency and the performance of the proposed approaches with respect to state-of-the-art solutions.

Contents

What is written without effort is in general read without pleasure.

Samuel Johnson

Introduction

Convex optimization is a rich and thriving mathematical discipline. While the theory of convex optimization has been largely explored for about a century, a new interest in the topic has been recently stimulated by two important developments. The first one is the emergence of efficient optimization methods, such as proximal algorithms, which allow one to easily solve large-size non-smooth convex problems. The second one is the discovery of the fact that many problems of practical interest can be tackled by non-smooth convex optimization.

In this thesis, we build a convex optimization framework to address three problems arising in image restoration, machine learning, and digital forensics. The common thread among these applications is the fact that the corresponding optimization problems involve nonlinear constraints difficult to handle with state-of-the-art solutions. Therefore, in the first part of the thesis, we develop a splitting technique to deal with a class of nonlinear convex constraints, whereas in the second part we focus on the specific contributions brought in the above applications. In the following, we summarize the content of the thesis.

Part I. Methodology

In the first part, we develop a technique to deal with a class of nonlinear convex constraints [50, 51]. The topic is organized in the following four chapters.

chapter 1 – convex optimization In this chapter, we recall the main tools provided by non-smooth convex optimization. We specifically focus on *proximal algorithms*, for which the key tool is the so-called *proximity operator*. We highlight the difficulty of proximal methods in addressing convex optimization problems involving "sophisticated" constraints. Indeed, when a hard constraint is involved, the proximity operator reverts to the *projection* onto the associated convex set, which is available in few cases only. We thus review the state-of-the-art approaches for computing the projection onto a convex set, and we motivate the interest of considering the splitting technique presented in Chapter 2.

chapter 2 – epigraphical splitting In this chapter, we deal with constraints expressed as the sublevel set of a block-separable function composed with a linear operator (e.g., the total variation penalty). Firstly, we present a splitting technique that replaces the aforementioned constraint with a collection of epigraphs. So doing, we trade the problem of computing the projection onto the former constraint set with the problem of computing the projection onto smaller epigraphs. Secondly, we enrich the list of functions for which the projec-

tion onto the associated epigraph can be efficiently computed. In this regard, we provide some theoretical results concerning the epigraphical projection of several functions of practical interest, such as the absolute value raised to a power $q \in]1, +\infty[$, the distance function to a convex set, the Euclidean norm, the infinity norm, and the max function.

CHAPTER 3 – CONSTRAINTS BASED ON MIXED NORMS In this chapter, we focus on constraints involving mixed norms, such as the total variation (TV) and nonlocal total variation (NLTV) penalties. We illustrate the performance of the epigraphical splitting technique through an example of image recovery from blurred, noisy, and decimated pixels. Numerical simulations demonstrate that constraints based on TV and NLTV measures can be efficiently handled by the epigraphical splitting, with significant improvements in terms of execution time w.r.t. standard numerical solutions.

CHAPTER 4 – CONSTRAINTS BASED ON PIECEWISE-AFFINE FUNCTIONS In this chapter, we deal with constraints expressed as the sublevel set of a sum of functions for which the associated epigraphical projection admits no closed-form expression. In such a case, we propose to outer approximate the original constraint by replacing each function involved in the sum with a piecewise-affine lower-approximating function. So doing, the epigraphical splitting decomposes the approximated constraint into a collection of convex polyhedrons. In order to handle the resulting constrained convex optimization problem, we present a primal-dual approach grounded on the epigraphical projection of the max function derived in Chapter 2. In addition, we demonstrate through numerical simulations that constraints based on Kullback-Leibler divergence and logistic loss can be efficiently handled by the proposed approximation method.

PART II. APPLICATIONS

In the second part, we present three problems arising in image restoration [53], supervised classification [52], and image forgery detection [49], which we propose to solve through the epigraphical splitting technique derived in Part I. The proposed approach is thoroughly discussed in the following three chapters.

CHAPTER 5 – NONLOCAL STRUCTURE TENSOR In this chapter, we extend the NLTV regularization to multicomponent images by taking advantage of the so-called *structure tensor* (ST). More specifically, we propose to penalize the nonlocal variations, jointly for the different components, through various mixed matrix norms. To facilitate the choice of the hyperparameters, we adopt a constrained formulation in which the data fidelity term is optimized over a constraint involving the ST-NLTV regularization. We solve the resulting convex optimization problem with the epigraphical splitting technique presented in Part I. Experiments carried out for color, multispectral and hyperspectral images demonstrate the interest of introducing a nonlocal structure tensor regularization, and show that the proposed epigraphical splitting leads to significant improvements in terms of execution time over current state-of-the-art methods.

CHAPTER 6 – SPARSE MULTICLASS SVM In this chapter, we propose a convex optimization approach for efficiently and exactly solving the multiclass SVM learning problem involving a sparse regularization and the multiclass hinge loss formulated by Crammer and Singer. We consider two approaches: one including the hinge loss as a penalty term, and the other one addressing the case when the hinge loss is enforced as a constraint. We implement these convex optimization problems through a primal-dual proximal algorithm and the epigraphical splitting technique presented in Part I. Experiments carried out for several datasets demonstrate the interest in considering the exact expression of the hinge loss rather than a smooth approximation, and show the efficiency of the proposed algorithms w.r.t. state-of-the-art methods.

CHAPTER 7 – IMAGE FORGERY DETECTION In this chapter, we propose a convex variational approach to detect image forgeries by using the photo response non-uniformity (PRNU), a deterministic pattern noise introduced by digital cameras in captured photos due to manufacturing imperfections of the imaging sensor. More specifically, we check the integrity of a photo by detecting the presence of the PRNU associated with its camera, dealing successfully with forgeries that elude most other detection strategies. Casting the problem in terms of Bayesian estimation, we use a suitable Markov random field prior for modeling the strong spatial dependencies of the source, and take decisions jointly on the whole image. This leads to a convex optimization problem involving a ℓ_1-norm constraint, which we efficiently solve through the epigraphical splitting technique presented in Part I. Large-scale experiments on simulated and real forgeries show that the proposed technique largely improves upon the current state-of-the-art, and that it can be applied with success to a wide range of practical situations.

Publications

This thesis has brought to the publication of the articles reported in the following.

Journal papers

(i). G. Chierchia, N. Pustelnik, J.-C. Pesquet, B. Pesquet-Popescu, "Epigraphical Projection and Proximal Tools for Solving Constrained Convex Optimization Problems", *Signal Image Video Process.*, July 2014.

(ii). G. Chierchia, G. Poggi, C. Sansone, L. Verdoliva, "A Bayesian-MRF Approach for PRNU-based Image Forgery Detection", *IEEE Trans. Inf. Forensics Security*, vol. 9, no. 4, pp. 554–567, Apr. 2014.

(iii). G. Chierchia, N. Pustelnik, B. Pesquet-Popescu, J.-C. Pesquet, "A Nonlocal Structure Tensor Approach for Multicomponent Image Recovery Problems", *IEEE Trans. Image Process.*, vol. 23, no. 12, pp. 5531–5544, Dec. 2014.

(iv). G. Chierchia, N. Pustelnik, J.-C. Pesquet, B. Pesquet-Popescu, "A Proximal Approach for Sparse Multiclass SVM", *J Mach Learn Res*, submitted, 2015.

(v). A. Fiengo, G. Chierchia, M. Cagnazzo and B. Pesquet-Popescu, "A convex optimization framework for optimal rate allocation in predictive video coding", *submitted*, 2015.

Conference papers

(i). G. Chierchia, N. Pustelnik, J.-C. Pesquet and B. Pesquet-Popescu, "A Proximal Approach for Constrained Cosparse Modelling", in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Kyoto, Japan, March 2012, pp. 3433–3436.

(ii). R. Gaetano, G. Chierchia and B. Pesquet-Popescu, "Non-local smoothness constraints for disparity estimation in a variational framework", in *European Signal Processing Conference*, Bucharest, Romania, Aug. 2012, pp. 1144–1148.

(iii). R. Gaetano, G. Chierchia and B. Pesquet-Popescu, "Parallel Implementations of a Disparity Estimation Algorithm based on a Proximal Splitting Method", in *Visual Communications and Image Processing*, San Diego, USA, Nov. 2012, pp. 1–6.

(iv). G. Chierchia, N. Pustelnik, J.-C. Pesquet and B. Pesquet-Popescu, "An Epigraphical Convex Optimization Approach for Multicomponent Image Restoration using Non-Local Structure Tensor", in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vancouver, Canada, May 2013, pp. 1359–1363.

(v). G. Chierchia, G. Poggi, C. Sansone and L. Verdoliva, "PRNU-based Forgery Detection with Regularity Constraints and Global Optimization", in *IEEE International Workshop on Multimedia Signal Processing*, Pula, Italy, Sep. 2013, pp. 236–241.

(vi). G. Chierchia, N. Pustelnik, J.-C. Pesquet and B. Pesquet-Popescu, "Epigraphical proximal projection for sparse Multiclass SVM", in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Florence, Italy, May 2014, pp. 8312–8316.

(vii). A. Fiengo, G. Chierchia, M. Cagnazzo and B. Pesquet-Popescu, "A convex-optimization framework for frame-level optimal rate allocation in predictive video coding", in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Florence, Italy, May 2014, pp. 7328–7332.

(viii). G. Chierchia, D. Cozzolino, G. Poggi, C. Sansone and L. Verdoliva, "Guided filtering for PRNU-based localization of small-size image forgeries", in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Florence, Italy, May 2014, pp. 6231–6235.

(ix). P. Lauga, G. Valenzise, G. Chierchia and F. Dufaux, "Improved Tone Mapping Operator for HDR Coding Optimizing the Distortion/Spatial Complexity Trade-Off", in *European Signal Processing Conference*, Lisbon, Portugal, June 2014, pp. 1607–1611.

(x). G. Chierchia, N. Pustelnik, J.-C. Pesquet and B. Pesquet-Popescu, "An Epigraphic Splitting Technique for Sparse Multiclass SVM", in *Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, Cambridge, UK, July 2015.

Ce qui est écrit sans effort
est généralement lu sans plaisir.

Samuel Johnson

Résumé

Dans cette thèse, nous proposons une approche d'optimisation convexe pour aborder des problèmes en restauration d'images multi-composantes, en apprentissage supervisé et en détection d'images falsifiées. Le fil conducteur de ces problèmes est la présence de contraintes convexes non linéaires qui sont difficiles à gérer avec les méthodes de l'état-de-l'art. Par conséquent, nous avons élaboré une technique d'éclatement épigraphique pour en simplifier la gestion. En s'appuyant sur cette approche, nous avons également proposé des contributions spécifiques pour les applications susmentionnées.

La première partie de la thèse décrit l'éclatement épigraphique de contraintes non-linéaires. Plus précisément, il s'agit de décomposer l'ensemble de sous-niveau d'une fonction séparable en une collection d'épigraphes. Ceci se traduit par une réduction de la complexité des algorithmes d'optimisation lorsque les épigraphes sont définis en termes de fonctions telles que la valeur absolue, la fonction distance, la norme euclidienne, la norme infinie, ou la fonction max. Nous démontrerons au travers de simulations numériques que la méthode proposée peut efficacement gérer des contraintes basées sur des fonctions répandues en restauration d'image et en apprentissage supervisé, comme la variation totale non-locale, la divergence de Kullback-Leibler ou la régression logistique.

La deuxième partie de la thèse apporte trois contributions fondées sur l'éclatement épigraphique. La première contribution concerne une nouvelle régularisation d'images multi-composants qui généralise la variation totale non-locale en prenant en compte le tenseur de structure. La seconde contribution fournit un algorithme d'apprentissage supervisé permettant d'apprendre efficacement et exactement une machine à vecteurs de support multi-classe avec une contrainte de parcimonie. Enfin, la troisième contribution porte sur une approche variationnelle pour la détection d'images falsifiées à travers la non-uniformité de la réponse photonique (bruit propre à chaque caméra). Au travers d'expériences numériques menées pour chaque application, nous démontrerons l'efficacité et la performance de notre approche par rapport aux méthodes de l'état-de-l'art.

Optimisation convexe

Les problèmes inverses peuvent rarement être résolus en manière exacte, à cause du fait qu'ils sont mal posés. L'objectif est donc de trouver une solution inexacte qui est optimale selon un critère statistique (ou dépendant de l'application), entrainant un problème d'optimisation comme dans la célèbre méthode des moindre carrées.

La résolution d'un problème d'optimisation est généralement basée sur des méthodes itératives qui calculent numériquement une solution du problème. Parmi ces méthodes, une grosse partie est consacrée aux problèmes convexes, qui sont le sujet centrale de cette thèse. En particulier, nous sommes intéressés à la résolution de problèmes d'optimisation convexe sous contrainte, dans la forme suivante:

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad f_0(x) \quad \text{s.t.} \quad (\forall r \in \{1,\dots,R\}) \quad f_r(x) \leq \eta_r.$$

L'intérêt dans ce type de problèmes vient du fait que les paramètres $(\eta_r)_{1\leq r\leq R}$ peuvent être facilement fixés s'ils sont en relation directe avec certains propriétés de la solution ou des données observées.

Un nouveau type de méthodes d'optimisation a été récemment proposé pour résoudre efficacement une large classe de problèmes convexes. Ces méthodes sont basées sur la notion d'opérateur proximal d'une fonction convexe semi-continue inférieurement $f \in \Gamma_0(\mathbb{R}^N)$:

$$(\forall x \in \mathbb{R}^N) \qquad \operatorname{prox}_f(x) = \underset{u\in\mathbb{R}^N}{\arg\min} \ \frac{1}{2}\|u-x\|_2^2 + f(u).$$

Quand f est la fonction indicatrice d'un ensemble convexe fermé non vide $C \subset \mathbb{R}^N$, c'est-à-dire

$$(\forall x \in \mathbb{R}^N) \qquad \iota_C(x) = \begin{cases} 0, & \text{if } x \in C, \\ +\infty, & \text{otherwise,} \end{cases}$$

l'opérateur proximal n'est rien d'autre que la projection orthogonale sur C, dans le sens que:

$$(\forall x \in \mathbb{R}^N) \qquad \operatorname{prox}_{\iota_C}(x) = P_C(x) = \underset{u\in C}{\arg\min} \ \|u-x\|_2^2.$$

Cela permet aux méthodes proximales de gérer soit des contraintes dures soit des fonctions non-lisses.

L'opérateur proximal constitue la brique de base de nombreux schémas d'éclatement tels que l'algorithme Forward-Backward ou les approches primaires-duales (deux exemples sont reportés dans les algorithmes 0.1 et 0.2). L'efficacité des méthodes d'éclatement

dépend donc de la capacité à calculer efficacement les opérateurs proximaux impliqués. Malheureusement, il existe de nombreux cas de projections pour lesquels la réponse est négative, ce qui explique l'intérêt de la méthode proposée dans cette thèse.

Algorithm 0.1 *Méthode FBPD*

INITIALIZATION

$$\left\lfloor \begin{array}{l} \text{choose } \left(x^{[0]}, y^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^M \\ \text{set } \tau > 0 \text{ and } \sigma > 0 \text{ such that} \\ \qquad \tau\left(\beta/2 + \sigma\|F\|^2\right) < 1 \end{array} \right.$$

FOR $i = 0, 1, \ldots$

$$\left\lfloor \begin{array}{l} \widehat{x}^{[i]} = \nabla g(x^{[i]}) - w + F^\top y^{[i]} \\ x^{[i+1]} = \operatorname{prox}_{\tau f}\left(x^{[i]} - \tau\, \widehat{x}^{[i]}\right) \\ \widehat{y}^{[i]} = F\big(2x^{[i+1]} - x^{[i]}\big) \\ y^{[i+1]} = \operatorname{prox}_{\sigma h^*}\left(y^{[i]} + \sigma\, \widehat{y}^{[i]}\right) \end{array} \right.$$

Algorithm 0.2 *Méthode M+LFBF*

INITIALIZATION

$$\left\lfloor \begin{array}{l} \text{choose } \left(x^{[0]}, y^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^M \\ \text{set } \gamma \in \left]0, (\beta + \|F\|)^{-1}\right[\end{array} \right.$$

FOR $i = 0, 1, \ldots$

$$\left\lfloor \begin{array}{l} \widehat{x}^{[i]} = \nabla g(x^{[i]}) - w + F^\top y^{[i]} \\ p^{[i]} = \operatorname{prox}_{\gamma f}\left(x^{[i]} - \gamma\, \widehat{x}^{[i]}\right) \\ v^{[i]} = \operatorname{prox}_{\gamma h^*}\left(y^{[i]} + \gamma\, F x^{[i]}\right) \\ y^{[i+1]} = v^{[i]} + \gamma F\left(p^{[i]} - x^{[i]}\right) \\ \widetilde{x}^{[i]} = \nabla g(p^{[i]}) - w + F^\top v^{[i]} \\ x^{[i+1]} = p^{[i]} - \gamma\left(\widetilde{x}^{[i]} - \widehat{x}^{[i]}\right) \end{array} \right.$$

Éclatement épigraphique

Plusieurs contraintes souvent impliquées dans la formulation de problèmes inverses sont basées sur une fonction décomposable, comme la divergence de Kullback-Leibler, la fonction « hinge » ou les normes mixtes. Nous proposons une nouvelle méthode pour traiter efficacement ce type de contraintes. En particulier, nous sommes intéressés à la classe suivante de problème d'optimisation convexe sous contrainte:

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\, x) \quad \text{s. t.} \quad \sum_{\ell=1}^{L} h_\ell(F_\ell\, x) \leq \eta,$$

où $\eta \in \mathbb{R}$ et

- $T_r \in \mathbb{R}^{K_r \times N}$ et $f_r \in \Gamma_0(\mathbb{R}^{K_r})$ pour chaque $r \in \{1, \dots, R\}$,
- $F_\ell \in \mathbb{R}^{M_\ell \times N}$ et $h_\ell \in \Gamma_0(\mathbb{R}^{M_\ell})$ pour chaque $\ell \in \{1, \dots, L\}$.

Notre idée consiste à introduire un vecteur auxiliaire $\zeta \in \mathbb{R}^L$ dans le problème d'optimisation, afin d'éclater la contrainte originale dans une collection d'épigraphes et un demi-espace fermé :

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^L}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\, x) \qquad \text{s. t.} \qquad \begin{cases} h_1\left(F_1\, x\right) \leq \zeta^{(1)} \\ \qquad \cdots \\ h_L\left(F_L\, x\right) \leq \zeta^{(L)} \\ \displaystyle\sum_{\ell=1}^{L} \zeta^{(\ell)} \leq \eta. \end{cases}$$

Le problème obtenu peut être efficacement résolu par les méthodes proximales de type primaire-duale, à condition que les projections épigraphiques $(P_{\operatorname{epi} h_\ell})_{1\leq\ell\leq L}$ soient rapidement calculables. Nous avons démontré que la projection épigraphique peut se ramener à l'opérateur proximale d'une fonction composée, comme illustré dans la proposition suivante.

Proposition 0.0.1. *Soit $\mathcal{H}$ un espace réel de Hilbert et soit $\mathcal{H} \times \mathbb{R}$ équipé avec le produit scalaire standard. Soit φ une fonction de $\Gamma_0(\mathcal{H})$ telle que* $\operatorname{dom}\varphi$ *est ouvert. Pour chaque $(y, \zeta) \in \mathcal{H} \times \mathbb{R}$, la projection dans* $\operatorname{epi}\varphi$ *est égale à*

$$P_{\operatorname{epi}\varphi}(y, \zeta) = \Big(p,\ \max\{\varphi(p), \zeta\}\Big),$$

où

$$p = \operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y).$$

A partir de ce résultat, nous avons dérivé une expression explicite de la projection épigraphique pour plusieurs fonctions d'intérêt, comme la valeur absolue, la fonction distance, la norme Euclidienne, la norme infinie et la fonction max. Ces nouvelles expressions sont présentées dans les propositions suivantes.

Proposition 0.0.2. *Soit* $\tau \in]0, +\infty[$ *et*

$$(\forall y \in \mathbb{R}) \qquad \varphi(y) = \tau|y|.$$

Pour chaque $(y, \zeta) \in \mathbb{R} \times \mathbb{R}$*, il résulte que*

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} y, & \text{if } \tau|y| \leq \zeta, \\ \dfrac{\operatorname{sign}(y)}{1+\tau^2} \max\{|y| + \tau\zeta, 0\}, & \text{otherwise.} \end{cases}$$

Proposition 0.0.3. *Soit* $q \in [1, +\infty[$, $\tau \in]0, +\infty[$ *et*

$$(\forall y \in \mathbb{R}) \qquad \varphi(y) = \tau|y|^q.$$

Pour chaque $(y, \zeta) \in \mathbb{R} \times \mathbb{R}$*, il résulte que*

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} y, & \text{if } \tau|y|^q \leq \zeta, \\ \operatorname{sign}(y)\, \chi, & \text{otherwise,} \end{cases}$$

où χ *est la solution unique sur* $[(\max\{\zeta, 0\}/\tau)^{1/q}, +\infty[$ *de l'équation*

$$q\tau^2\chi^{2q-1} - q\tau\zeta\chi^{q-1} + \chi = |y|.$$

Proposition 0.0.4. *Soit* C *un sous-ensemble convexe (non vide) de* $\mathcal{H}$*. Soit* $q \in [1, +\infty[$, $\tau \in]0, +\infty[$, $\zeta \in \mathbb{R}$ *et*

$$(\forall y \in \mathcal{H}) \qquad \varphi(y) = \tau d_C^q(y).$$

Pour chaque $(y, \zeta) \in \mathcal{H} \times \mathbb{R}$*, il résulte que*

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} y, & \text{if } y \in C, \\ \alpha y + (1-\alpha) P_C(y), & \text{otherwise,} \end{cases}$$

où

$$\alpha = \frac{\operatorname{prox}_{\frac{1}{2}(\max\{\tau|\cdot|^q-\zeta,0\})^2}\left(d_C(y)\right)}{d_C(y)}.$$

Corollary 0.0.5. *Soit* $\tau \in]0, +\infty[$, $\zeta \in \mathbb{R}$, $z \in \mathcal{H}$ *et*

$$(\forall y \in \mathcal{H}) \qquad \varphi(y) = \tau \|y - z\|_2.$$

Pour chaque $(y, \zeta) \in \mathcal{H} \times \mathbb{R}$, *il résulte que*

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} z, & \text{if } y = z, \\ y, & \text{if } \tau\|y - z\|_2 \leq \zeta, \\ z + \alpha(y - z), & \text{otherwise}, \end{cases}$$

où

$$\alpha = \frac{1}{1 + \tau^2} \max\left\{1 + \frac{\tau\zeta}{\|y - z\|_2}, 0\right\}.$$

Les derniers résultats concernent la projection épigraphique de la fonction max et de la norme infini. La preuve de ces propositions est basée sur l'opérateur proximale introduit dans le lemme suivant.

Lemma 0.0.6. *Soit* $(\tau_m)_{1\leq m\leq M} \in \mathbb{R}^M$, $\nu = (\nu^{(m)})_{1\leq m\leq M} \in \mathbb{R}^M$ *et*

$$(\forall y \in \mathbb{R}) \qquad \varphi(y) = \frac{1}{2} \sum_{m=1}^{M} \Big(\max\{\tau_m \, (\nu^{(m)} - y), 0\}\Big)^2,$$

avec les valeurs $(\nu^{(m)})_{1\leq m\leq M_\ell}$ *triés en ordre croissant. Il résulte que* $\varphi \in \Gamma_0(\mathbb{R})$ *et*

$$(\forall y \in \mathbb{R}) \qquad \operatorname{prox}_{\varphi}(y) = \frac{y + \sum_{m=1}^{\overline{m}-1} \nu^{(m)}(\tau_m^-)^2 + \sum_{m=\overline{m}}^{M} \nu^{(m)}(\tau_m^+)^2}{1 + \sum_{m=1}^{\overline{m}-1} (\tau_m^-)^2 + \sum_{m=\overline{m}}^{M} (\tau_m^+)^2},$$

où $\tau_m^- = \min\{\tau_m, 0\}$ *et* $\tau_m^+ = \max\{\tau_m, 0\}$ *pour chaque* $m \in \{1, \ldots, M\}$, *alors que* $\overline{m}$ *est l'unique entier dans* $\{1, \ldots, M + 1\}$ *tel que*

$$\nu^{(\overline{m}-1)} < \frac{y + \sum_{m=1}^{\overline{m}-1} \nu^{(m)}(\tau_m^-)^2 + \sum_{m=\overline{m}}^{M} \nu^{(m)}(\tau_m^+)^2}{1 + \sum_{m=1}^{\overline{m}-1} (\tau_m^-)^2 + \sum_{m=\overline{m}}^{M} (\tau_m^+)^2} \leq \nu^{(\overline{m})},$$

avec la convention $\nu^{(0)} = -\infty$, $\nu^{(M+1)} = +\infty$ *et* $\sum_{m=1}^{0} \cdot = \sum_{m=M+1}^{M} \cdot = 0$.

Proposition 0.0.7. *Soit* $(\tau_m)_{1\leq m\leq M} \in]0,+\infty[^M$ *et*

$$(\forall y \in \mathbb{R}^M) \qquad \varphi(y) = \max_{1\leq m\leq M} \tau_m \, |y^{(m)}|,$$

avec les valeurs $\left(\nu^{(m)} = \tau_m \, |y^{(m)}|\right)_{1\leq m\leq M}$ *triés en ordre croissant. Pour chaque* $(y,\zeta) \in \mathbb{R}^M \times \mathbb{R}$*, la projection* $P_{\mathrm{epi}\,\varphi}(y,\zeta) = (p,\theta)$ *est*

$$p = \left[\operatorname{sign}(y^{(m)}) \, \min\left\{ |y^{(m)}|, \theta/\tau_m \right\} \right]_{1\leq m\leq M}$$

et

$$\theta = \max\left\{ 0, \, \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^{-2}\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^{-2} \, \nu^{(m)}\right) \right\},$$

où $\overline{m}$ *est l'unique entier dans* $\{1,\dots,M+1\}$ *tel que*

$$\nu^{(\overline{m}-1)} < \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^{-2}\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^{-2} \, \nu^{(m)}\right) \leq \nu^{(\overline{m})},$$

avec les conventions $\nu^{(0)} = -\infty$*,* $\nu^{(M+1)} = +\infty$ *et* $\sum_{m=M+1}^{M} \cdot = 0$*.*

Proposition 0.0.8. *Soit* $(r_m)_{1\leq m\leq M} \in \mathbb{R}^M$*,* $(\tau_m)_{1\leq m\leq M} \in]0,+\infty[^M$*,* $(\epsilon_m)_{1\leq m\leq M} \in \{1,-1\}^M$ *et*

$$(\forall y \in \mathbb{R}^M) \qquad \varphi(y) = \max_{1\leq m\leq M} \frac{\epsilon_m \, y^{(m)} + r_m}{\tau_m},$$

avec les valeurs $\left(\nu^{(m)} = \frac{\epsilon_m y^{(m)} + r_m}{\tau_m}\right)_{1\leq m\leq M}$ *triés en ordre croissant. Pour chaque* $(y,\zeta) \in \mathbb{R}^M \times \mathbb{R}$*, la projection* $P_{\mathrm{epi}\,\varphi}(y,\zeta) = (p,\theta)$ *est*

$$p = \left[\epsilon_m \, \min\left\{ \epsilon_m \, y^{(m)}, \tau_m \, \theta - r_m \right\} \right]_{1\leq m\leq M},$$

et

$$\theta = \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^{2}\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^{2} \, \nu^{(m)}\right),$$

où $\overline{m}$ *est l'unique entier dans* $\{1,\dots,M+1\}$ *tel que*

$$\nu^{(\overline{m}-1)} < \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^{2}\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^{2} \, \nu^{(m)}\right) \leq \nu^{(\overline{m})},$$

avec les conventions $\nu^{(0)} = -\infty$*,* $\nu^{(M+1)} = +\infty$ *et* $\sum_{m=M+1}^{M} \cdot = 0$*.*

Contraintes basées sur les normes mixtes

L'éclatement épigraphique trouve une application immédiate aux contraintes basées sur les normes mixtes, lesquelles sont souvent utilisées en restauration d'images. Dans ce scénario, l'objectif est de retrouver une image $\overline{x} \in \mathbb{R}^N$ à partir d'une observation $z \in \mathbb{R}^K$ bruitée par un modèle linéaire

$$z = A\overline{x} + b,$$

où $A \in \mathbb{R}^{K \times N}$ (avec $K \leq N$) est un opérateur connu, et $b \in \mathbb{R}^K$ est une réalisation de bruit Gaussien à moyenne nulle.

Une approche usuelle pour estimer $\overline{x}$ à partir de z consiste à résoudre un problème d'optimisation convexe du type:

$$\underset{x \in [0,255]^N}{\text{minimize}} \quad \|Ax - z\|_2^2 \quad \text{s. t.} \quad \sum_{\ell=1}^{N} \|F_\ell\, x\|_p \leq \eta,$$

où $\eta \geq 0$ et $p \geq 1$. Pour la contrainte de régularisation, nous utilisons la variation totale non-locale, qui corresponde à choisir un F_ℓ tel que

$$F_\ell\, x = \left[\omega_{\ell,n}(x^{(\ell)} - x^{(n)})\right]_{n \in \mathcal{N}_\ell},$$

où $\mathcal{N}_\ell \subset \{1, \ldots, N\} \setminus \{\ell\}$ contient des positions autour de ℓ, et $(\omega_{\ell,n})_{n \in \mathcal{N}_\ell}$ sont des poids positifs qui mesurent la similarité entre $x^{(\ell)}$ et ses voisins $(x^{(n)})_{n \in \mathcal{N}_\ell}$.

L'éclatement épigraphique permet de résoudre le problème susmentionné sans boucles internes (contrairement aux approches standards), grâce au fait que la contrainte originale est décomposée en épigraphes de la norme $\boldsymbol{\ell}_p$ (pour lesquelles nous avons dérivé une expression explicite de la projection) :

$$\underset{(x,\zeta) \in [0,255]^N \times \mathbb{R}^N}{\text{minimize}} \quad \|Ax - z\|_2^2 \quad \text{s. t.} \quad \begin{cases} \left(F_1\, x, \zeta^{(1)}\right) \in \operatorname{epi} \|\cdot\|_p \\ \qquad \cdots \\ \left(F_L\, x, \zeta^{(L)}\right) \in \operatorname{epi} \|\cdot\|_p \\ \sum_{\ell=1}^{L} \zeta^{(\ell)} \leq \eta. \end{cases}$$

Cela permet d'obtenir un algorithme rapide qui peut traiter efficacement des images multi-canal dégradées par un opérateur de flou, un bruit additif gaussien, et une perte d'une partie des pixels, comme illustré dans les figures 1-2. Également, la figure 3 montre que l'algorithme épigraphique est plus rapide (en temps d'exécution) que celui basé sur l'approche standard.

(a) Image originelle (zoom).

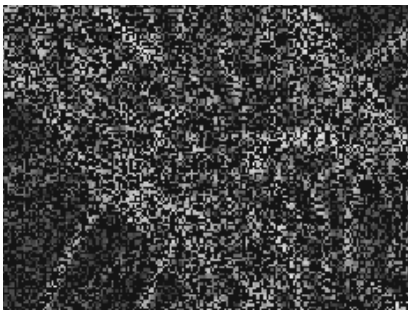

(b) Degradée.

(c) TV_1: 19.79–0.838.

(d) TV_2: 20.80–0.855.

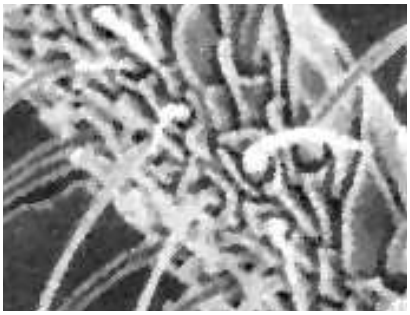

(e) TV_∞: 20.25–0.853.

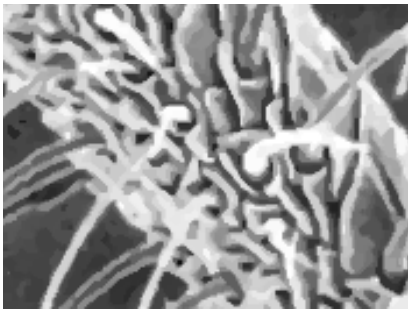

(f) $\mathrm{NLTV}_{1,1}$: 20.93–0.865.

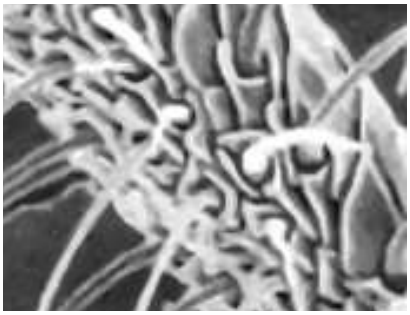

(g) $\mathrm{NLTV}_{1,2}$: **22.62–0.897**.

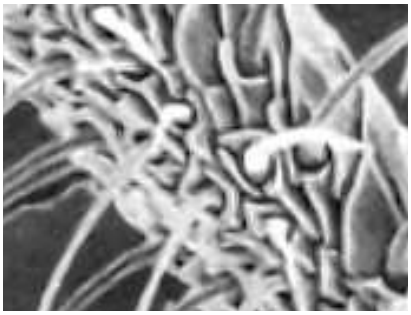

(h) $\mathrm{NLTV}_{1,\infty}$: 22.38–0.897.

Figure 1 *SNR (dB) et SSIM d'une image (en niveaux de gris) dégradée avec un flou uniforme de taille* 3×3*, un bruit de variance* 10^2*, et* 60% *de décimation.*

(a) Image originelle (zoom).

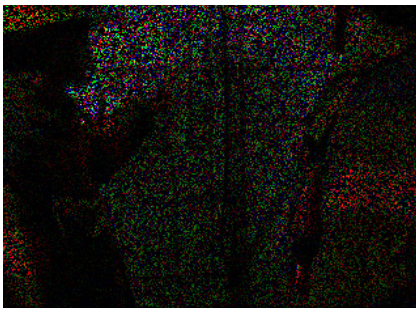

(b) Degradée.

(c) TV_1: 17.78–0.787.

(d) TV_2: 18.36–0.821.

(e) TV_∞: 18.91–0.824.

(f) $\mathrm{NLTV}_{1,1}$: 18.93–0.828.

(g) $\mathrm{NLTV}_{1,2}$: 19.47–0.839.

(h) $\mathrm{NLTV}_{1,\infty}$: **20.17–0.847**.

Figure 2 *SNR (dB) et SSIM d'une image (en couleurs) dégradée avec un flou uniforme de taille* 3×3*, un bruit de variance* 10^2*, et* 60% *de décimation.*

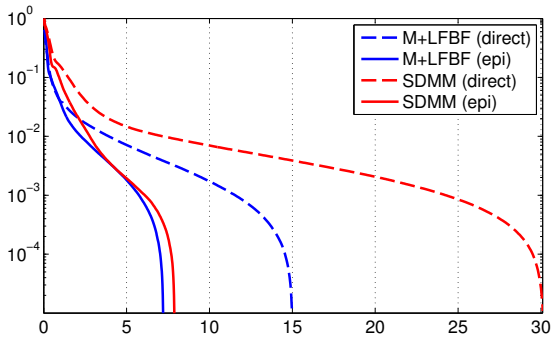

(a) TV_2.

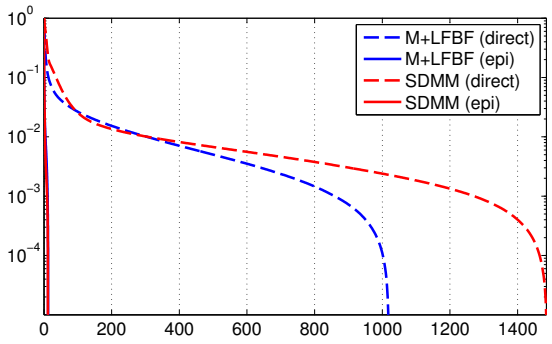

(b) TV_∞.

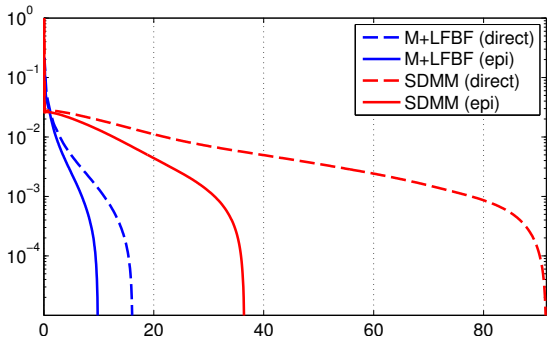

(c) $\mathrm{NLTV}_{1,2}$.

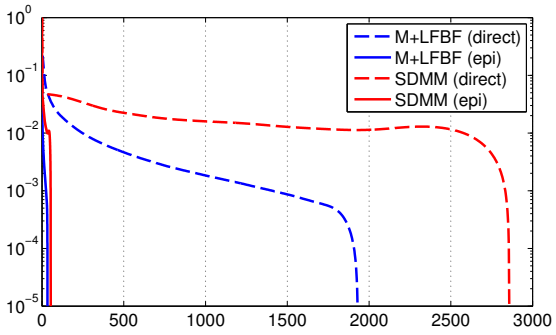

(d) $\mathrm{NLTV}_{1,\infty}$.

Figure 3 *Distance relative de $x^{[\infty]}$ vs temps d'exécution (en seconds): plots qui comparent les approches épigraphique et standard par rapport à deux méthodes proximales (M+LFBF et SDMM).*

Contraintes basées sur les fonctions affines

L'éclatement épigraphique est basé sur l'idée de remplacer la contrainte sur une fonction séparable avec une collection d'épigraphes:

$$\sum_{\ell=1}^{L} h_\ell(F_\ell\, x) \leq \eta \quad \Leftrightarrow \quad \begin{cases} (\forall \ell \in \{1,\ldots,L\}) \quad (F_\ell\, x, \zeta^{(\ell)}) \in \operatorname{epi} h_\ell, \\ \zeta^{(1)} + \cdots + \zeta^{(L)} \leq \eta. \end{cases}$$

Cette décomposition s'avère très efficace lorsque les fonctions $(h_\ell)_{1\leq\ell\leq L}$ sont des normes $\boldsymbol{\ell}_p$ avec $p \in \{1, 2, +\infty\}$, parce que la projection épigraphique correspondante est calculable rapidement. Cependant, cela n'est pas nécessairement vrai pour des fonctions h_ℓ qui n'ont pas de projection épigraphique en forme explicite.

Pour surmonter la difficulté susmentionnée, nous proposons d'approximer les fonctions $(h_\ell)_{1\leq\ell\leq L}$ impliquées dans la contrainte originale par des fonctions affines telles que

$$(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad \widehat{h}_\ell(\mathsf{y}^{(\ell)}) = \max_{1\leq j\leq J_\ell} \delta_{\ell,j}^\top\, \mathsf{y}^{(\ell)} + \mu_{\ell,j},$$

où les paires $(\delta_{\ell,j}, \mu_{\ell,j}) \in \mathbb{R}^{M_\ell} \times \mathbb{R}$ sont définies à partir d'un ensemble de valeurs distinctes $\{\mathsf{a}_{\ell,1}, \ldots, \mathsf{a}_{\ell,J_\ell}\} \subset \mathbb{R}^{M_\ell}$ de la manière suivante :

$$\delta_{\ell,j} \in \partial h_\ell(\mathsf{a}_{\ell,j}), \qquad \mu_{\ell,j} = h_\ell(\mathsf{a}_{\ell,j}) - \delta_{\ell,j}^\top \mathsf{a}_{\ell,j}.$$

La figure 4 montre une fonction et son approximation affine. Bien évidemment, la qualité de l'approximation peut être contrôlée par le nombre J_ℓ de paires $\{(\delta_{\ell,j}, \mu_{\ell,j})\}_{1\leq j\leq J_\ell}$ utilisées dans la fonction $\widehat{h}_\ell$.

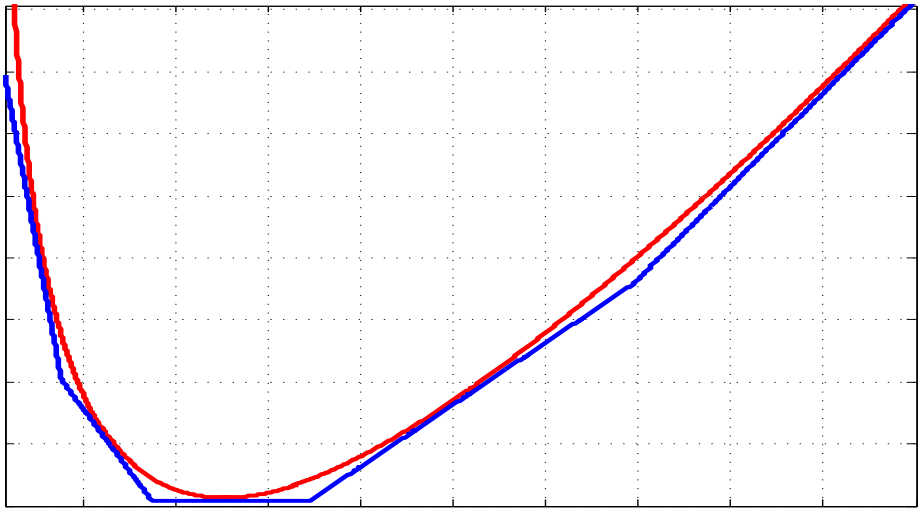

Figure 4 *Une fonction (trait rouge) et son approximation affine (trait bleu).*

L'éclatement épigraphique permet de résoudre efficacement le problème d'optimisation « approximé ». En fait, la fonction affine $\widehat{h}_\ell$ peut s'exprimer comme le max composée par un opérateur linéaire:

$$(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad \widehat{h}_\ell(\mathsf{y}^{(\ell)}) = h_\ell^{\max}(\Delta_\ell \, \mathsf{y}^{(\ell)}),$$

où $\Delta_\ell = [\delta_{\ell,1} \dots \delta_{\ell,J_\ell}]^\top$ et $h_\ell^{\max}(\mathsf{u}^{(\ell)}) = \max_{1 \le j \le J_\ell} \; u^{(\ell,j)} + \mu_{\ell,j}$. Ainsi, la décomposition épigraphique du problème approximé est

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^L}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\, x) \quad \text{s. t.} \quad \begin{cases} (\forall \ell \in \{1,\dots,L\}) \quad h_\ell^{\max}\left(\Delta_\ell F_\ell \, x\right) \le \zeta^{(\ell)}, \\ \displaystyle\sum_{\ell=1}^{L} \zeta^{(\ell)} \le \eta. \end{cases}$$

En fixant $F = [F_1^\top \; \dots \; F_L^\top]^\top \in \mathbb{R}^{M\times N}$ et $\Delta = \operatorname{diag}(\Delta_1, \dots, \Delta_L) \in \mathbb{R}^{J\times M}$ (avec $M = M_1 + \dots + M_L$ and $J = J_1 + \dots + J_L$), nous obtenons une forme compacte du problème précédent :

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^L}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\, x) \quad \text{s. t.} \quad \begin{cases} (\Delta F x, \zeta) \in E, \\ \zeta \in V, \end{cases}$$

où

$$E = \Big\{(u,\zeta) \in \mathbb{R}^J \times \mathbb{R}^L \;\Big|\; (\forall \ell \in \{1,\dots,L\}) \quad (\mathsf{u}^{(\ell)}, \zeta^{(\ell)}) \in \operatorname{epi} h_\ell^{\max}\Big\},$$
$$V = \Big\{\zeta \in \mathbb{R}^L \;\Big|\; 1_L^\top \zeta \le \eta\Big\}.$$

L'avantage de cette reformulation est que la projection dans E équivaut à calculer des projections épigraphiques de la fonction max, pour laquelle nous avons dérivé une forme explicite.

Une possible application de cette approche est la restauration d'une image $\overline{x} \in \mathbb{R}^N$ dégradé avec un opérateur de flou $A \in \mathbb{R}^{K\times N}$ et un bruit de Poission, c'est-à-dire $z = \mathcal{P}_\alpha(A\overline{x})$. Une possible approche pour récupérer $\overline{x}$ à partir de z consiste à formuler un problème d'optimisation convexe dans lequel une régularisation basée sur la variation totale (TV) est optimisée sous la contrainte de la divergence de Kullback-Leibler, à savoir

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad \mathrm{TV}_2(x) \quad \text{s. t.} \quad \sum_{\ell=1}^{K} h_\ell^{\mathrm{KL}}\Big((Ax)^{(\ell)}\Big) \le K/2.$$

Ce problème peut être approximé par la formulation présentée en haut. La Figure 5 montre une image dégradée par un flou uniforme de taille 3×3 et un bruit de Poission, ainsi que les images restaurées par la méthode proposée et l'état de l'art. Les figures indiquent que les résultats obtenus par la formulation approximée sont très proches de la solution « exacte », avec des temps de calcul acceptables.

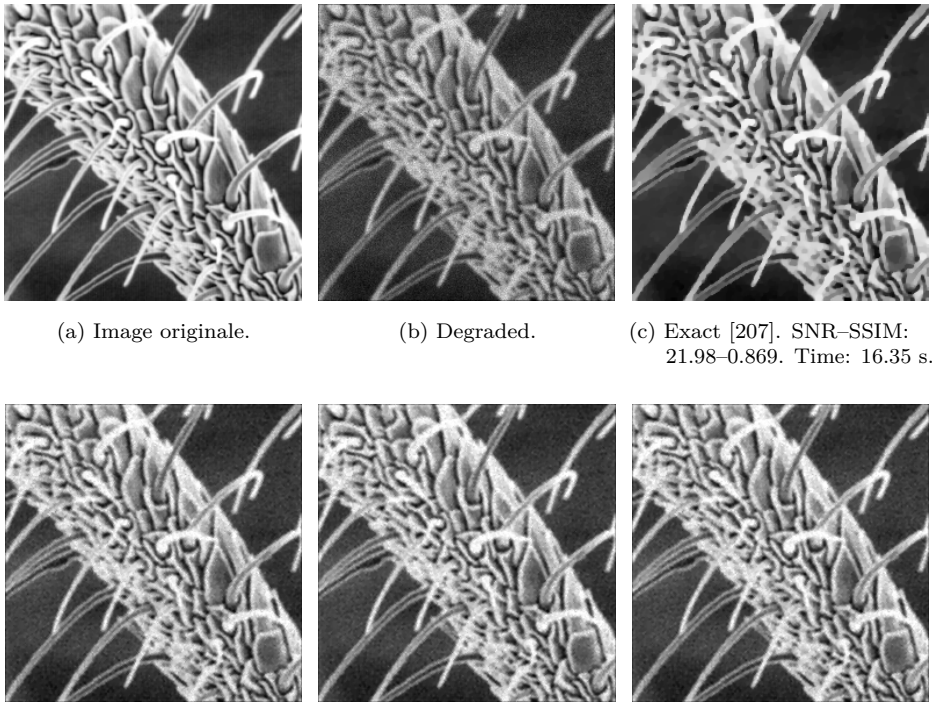

(a) Image originale. (b) Degraded. (c) Exact [207]. SNR–SSIM: 21.98–0.869. Time: 16.35 s.

(d) Approx with $J_k = 9$. SNR–SSIM: 21.80–0.861. Time: 44.05 s. (e) Approx with $J_\ell = 13$. SNR–SSIM: **22.04–0.862**. Time: 59.47 s. (f) Approx with $J_\ell = 17$. SNR–SSIM: 22.04–0.862. Time: 70.25 s.

Figure 5 *SNR (dB) – SSIM indexes for the recovery of a grayscale image degraded with a uniform blur of size* 3×3 *and a Poisson noise with scaling* $\alpha = 1$.

Une deuxième application de cette approche est utilisée pour faire de la classification parcimonieuse par régression logistique, ce qui consiste à résoudre le problème suivant:

$$\underset{x \in \mathbb{R}^{MK}}{\text{minimize}} \quad \sum_{k=1}^{K} \|x^{(k)}\|_{1,p} \quad \text{s. t.} \quad \sum_{\ell=1}^{L} h_\ell^{\log}(F_\ell\, x) \leq \eta,$$

où $h_\ell^{\log}(\mathsf{y}^{(\ell)}) = \log\left(\sum_{k=1}^{K} \exp\left(y^{(\ell,k)}\right)\right)$. Ce problème peut être également approximé par la formulation présentée en haut. La figure 6 montre un jeu de données contenant un large nombre de 28×28 images ($N = 784$) de chiffres manuscrites ($K = 10$) organisées en 60000 images de « training » et 10000 images de « test ». Le tableau 1 montre les erreurs de classification obtenues en faisant l'apprentissage sur un sous-ensemble de ces données de taille $L \in \{5K, 10K, 20K, 50K\}$ et avec plusieurs valeurs de η. Les résultats indiquent que l'approche proposée n'est que légèrement sensitive à l'approximation introduite.

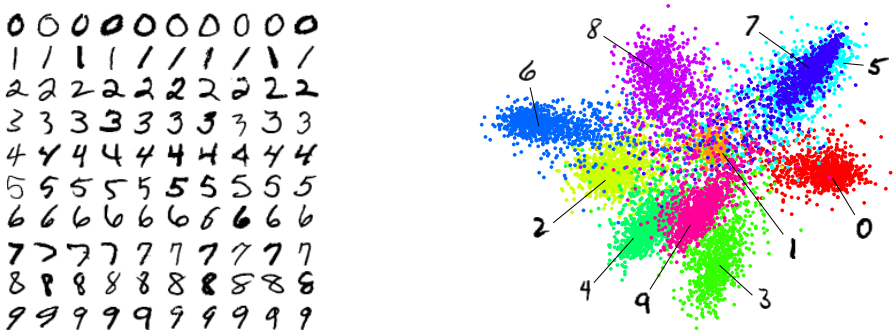

(a) Base de données. (b) Espace des carachteristiques.

Figure 6 *Database MNIST.*

Table 1 *Erreurs de classification (dans le cas $K = 3$) obtenus avec la regression logistique approximée pour plusieurs valeurs de L, η et $J_\ell = (2m+1)^{K-1}$.*

	η	$g = \|\| \cdot \|\|_{1,\infty}$				
		$m = 3$	$m = 5$	$m = 7$	$m = 9$	exact
$L/K = 5$	$0.01L$	3.37 %	3.40 %	3.37 %	3.43 %	2.76 %
	$0.05L$	3.08 %	3.53 %	3.24 %	3.21 %	
	$0.1L$	3.11 %	2.83 %	3.08 %	2.92 %	
	$0.5L$	2.45 %	3.08 %	3.11 %	**2.57** %	
	$0.7L$	3.15 %	2.64 %	2.80 %	2.80 %	
	$1.0L$	6.04 %	6.32 %	6.70 %	6.93 %	
$L/K = 10$	$0.01L$	3.05 %	2.96 %	2.89 %	2.92 %	2.51 %
	$0.05L$	2.48 %	2.48 %	2.86 %	2.48 %	
	$0.1L$	2.67 %	2.64 %	2.80 %	2.83 %	
	$0.5L$	2.32 %	2.67 %	2.51 %	**2.29** %	
	$0.7L$	2.80 %	2.70 %	2.41 %	2.41 %	
	$1.0L$	6.77 %	4.26 %	3.62 %	3.56 %	
$L/K = 20$	$0.005L$	1.78 %	2.22 %	2.03 %	2.07 %	1.97 %
	$0.01L$	2.35 %	2.07 %	2.00 %	2.07 %	
	$0.05L$	1.68 %	2.03 %	1.97 %	2.05 %	
	$0.1L$	2.07 %	2.07 %	1.94 %	**2.03** %	
	$0.5L$	2.32 %	2.76 %	2.67 %	2.70 %	
	$0.7L$	3.53 %	2.99 %	3.02 %	2.92 %	
	$1.0L$	8.64 %	6.23 %	7.05 %	6.55 %	
$L/K = 50$	$0.005L$	1.56 %	1.08 %	1.08 %	**1.14** %	1.49 %
	$0.01L$	1.56 %	1.43 %	1.18 %	1.53 %	
	$0.05L$	1.84 %	1.43 %	1.53 %	1.56 %	
	$0.1L$	2.03 %	1.84 %	1.81 %	1.84 %	
	$0.5L$	2.89 %	2.86 %	2.76 %	2.80 %	
	$0.7L$	4.26 %	3.37 %	3.37 %	3.34 %	
	$1.0L$	12.81 %	10.84 %	12.20 %	11.63 %	

Tenseur de structure non-local

Les images multi-composants sont un ensemble de cartes spatiales acquises simultanément, telles que les images couleurs (qui sont constituées par les canaux rouge, vert et bleu) ou les images spectrales (qui décomposent le spectre électromagnétique en plusieurs intervalles de fréquence). En raison d'imprécisions du senseur ou de limitations physiques, une image multi-composants $\overline{\mathrm{x}} = (\overline{x}_1, \ldots, \overline{x}_R) \in (\mathbb{R}^N)^R$ est souvent dégradées par un opérateur linéaire $A\colon (\mathbb{R}^N)^R \mapsto (\mathbb{R}^K)^S$ et un bruit $\mathcal{B}$, entrainant un modèle du type

$$\mathrm{z} = \mathcal{B}(\mathrm{A}\overline{\mathrm{x}}),$$

lequel se retrouve en plusieurs applications, telles que l'acquisition comprimée, la super-résolution ou le démélange spectrale.

L'approche usuelle pour estimer $\overline{x}$ à partir de z consiste à résoudre un problème d'optimisation convexe du type

$$\underset{\mathrm{x}\in C}{\text{minimize}} \quad f(\mathrm{Ax}, \mathrm{z}) \quad \text{s. t.} \quad h(\mathrm{x}) \leq \eta.$$

Le terme d'attache aux données $f(\cdot, \mathrm{z})$ permet d'assurer que la solution soit proche à l'image observée, dont l'expression dépende de la distribution statistique du bruit, pouvant être une fonction quadratique dans le cas Gaussien, une norme $\boldsymbol{\ell}_1$ dans le cas Laplacien, ou une divergence dans le cas Poissonien. En revanche, la fonction h permet d'imposer une contrainte de régularité sur la solution.

Nous proposons de modéliser les dépendances spatiales et spectrales d'une image multi-composants par une régularisation fondée sur des normes matricielles, à savoir

$$\left(\forall \mathrm{x} \in (\mathbb{R}^N)^R\right) \qquad h(\mathrm{x}) = \sum_{\ell=1}^{N} \|\mathrm{X}^{(\ell)}\|_p.$$

Ci-dessus, le symbole $\|\cdot\|_p$ dénote la norme de Schatten avec $p \geq 1$, tandis que $\mathrm{X}^{(\ell)}$ représente la matrice de gradients non-locales

$$\mathrm{X}^{(\ell)} = \left[\omega_{\ell,n}(x_r^{(\ell)} - x_r^{(n)})\right]_{n\in\mathcal{N}_\ell, 1\leq r\leq R},$$

où $\mathcal{N}_\ell \subset \mathcal{W}_\ell \setminus \{\ell\}$ est le support d'un voisinage de ℓ de taille M_ℓ. Par conséquent, le problème d'optimisation évoqué en haut peut être reformulé de la manière suivante

$$\underset{\mathrm{x}\in C}{\text{minimize}} \quad f(\mathrm{Ax}, \mathrm{z}) \quad \text{s. t.} \quad \Phi\,\mathrm{x} \in D,$$

où $M = M_1 + \cdots + M_N$, $\Phi\colon (\mathbb{R}^N)^R \mapsto \mathbb{R}^{M\times R}$ est l'opérateur linéaire

$$\Big(\forall \mathrm{x} \in (\mathbb{R}^N)^R\Big) \qquad \Phi \mathrm{x} = \mathrm{X} = \begin{bmatrix} \mathrm{X}^{(1)} \\ \dots \\ X^{(N)} \end{bmatrix}$$

et D est l'ensemble convexe défini comme

$$D = \Big\{ \mathrm{X} \in \mathbb{R}^{M\times R} \;\Big|\; \sum_{\ell=1}^{N} \|\mathrm{X}^{(\ell)}\|_p \leq \eta \Big\}.$$

L'approche épigraphique permet de résoudre efficacement ce problème en décomposant la contrainte D dans une collection d'épigraphes de la norme $\boldsymbol{\ell}_p$ appliquée aux valeur singulières de $\mathrm{X}^{(\ell)}$, dénotées dans la suite par $\sigma_{\mathrm{X}^{(\ell)}} = \Big(\sigma^{(m)}_{\mathrm{X}^{(\ell)}}\Big)_{1\leq m\leq \widetilde{M}_\ell}$, avec $\widetilde{M}_\ell = \min\{M_\ell, R\}$.

- Normes nucléaires (cas $p = 1$)

$$\begin{cases} (\forall \ell \in \{1,\dots,N\})(\forall m \in \{1,\dots,\widetilde{M}_\ell\}) \quad \Big|\sigma^{(m)}_{\mathrm{X}^{(\ell)}}\Big| \leq \zeta^{(\ell,m)}, \\ \displaystyle\sum_{\ell=1}^{N}\sum_{m=1}^{\widetilde{M}_\ell} \zeta^{(\ell,m)} \leq \eta. \end{cases}$$

- Normes de Frobenius ou spectrales (cas $p \in \{2, +\infty\}$)

$$\begin{cases} (\forall \ell \in \{1,\dots,N\}) \quad \|\sigma_{\mathrm{X}^{(\ell)}}\|_p \leq \zeta^{(\ell)}, \\ \displaystyle\sum_{\ell=1}^{N} \zeta^{(\ell)} \leq \eta. \end{cases}$$

L'avantage de cette approche est que la projection épigraphique d'une norme matricielle avec $p \in \{1, 2, +\infty\}$ peut être calculée en forme explicite. Les expériences numériques et les comparaisons avec l'état de l'art (notamment les méthodes ADMM) sont en faveur de l'approche proposée en terme de PSNR-SSIM et en temps de calcul. Les résultats en Fig. 7 et 8, ainsi que Tab. 2, montrent les meilleures performances de la régularisation basée sur le tenseur de structure par rapport à plusieurs types d'image multi-composant.

Les résultats montrent également que la norme nucléaire est plus indiquées que les autre normes matricielles pour traiter des images hyper-spectrales. De plus, l'algorithme épigraphique converges plus rapidement que les approches classiques basées sur le calcul itératif des projections à chaque itération. Dans tout les cas, l'algorithme proposé s'avère plus efficace que les solutions basées sur ADMM, indiquant que les méthodes proximales primaires-duales constituent un bon choix pour attaquer des problèmes de restauration d'images multi-composants.

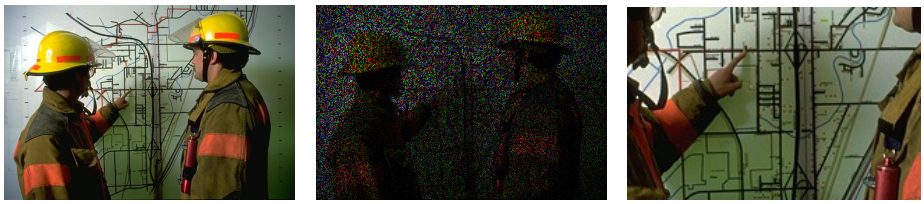

(a) Originale. (b) Bruitée. (c) Zoom.

(d) ℓ_1-CC-TV [6]: 16.15 dB (e) ℓ_2-CC-TV [6]: 16.32 dB (f) ℓ_∞-CC-TV [6]: 16.05 dB

(g) ℓ_1-CC-NLTV[51]: 16.87 dB (h) ℓ_2-CC-NLTV[51]: 17.20 dB (i) ℓ_∞-CC-NLTV[51]: 17.22 dB

(j) ℓ_1-ST-TV: 17.08 dB (k) ℓ_2-ST-TV [28]: 16.84 dB (l) ℓ_∞-ST-TV [108]: 16.43 dB

(m) ℓ_1-ST-NLTV: **18.20 dB** (n) ℓ_2-ST-NLTV: 17.46 dB (o) ℓ_∞-ST-NLTV: 16.67 dB

Figure 7 *Comparaison visuelle d'une image couleur reconstruite avec plusieurs contraintes de régularisation. Dégradation: bruit Gaussien additive avec moyenne nulle et écart type égal à* 10*, flou uniforme de taille* 3×3*, et* 80% *de décimation (*$N = 154401$*,* $R = S = 3$ *et* $K = 30880$*).*

Table 2 *Indices de SNR – M-SNR résultants de la reconstruction d'images dégradées par un bruit Gaussien avec écart type égal à* 5 *et* 90% *de décimation.*

image	size	H-TV [235]	ℓ_1-ST-TV	M-NLTV [45]	ℓ_1-ST-NLTV
Hydice	$256 \times 256 \times 191$	10.65 – 09.87	11.93 – 11.16	11.57 – 10.76	**12.98 – 12.11**
Indian Pine	$145 \times 145 \times 200$	17.31 – 17.00	18.46 – 18.24	17.62 – 17.34	**19.53 – 19.49**
Little River	$512 \times 512 \times 7$	17.81 – 18.20	18.49 – 18.83	18.46 – 18.90	**19.88 – 20.18**
Mississippi	$512 \times 512 \times 7$	18.27 – 18.07	18.60 – 18.37	18.94 – 18.59	**19.56 – 19.28**
Montana	$512 \times 512 \times 7$	22.49 – 20.97	22.68 – 21.15	22.85 – 21.29	**23.31 – 21.76**
Rio	$512 \times 512 \times 7$	16.48 – 15.29	16.65 – 15.48	16.82 – 15.64	**17.20 – 16.05**
Paris	$512 \times 512 \times 7$	14.85 – 14.31	14.94 – 14.39	15.05 – 14.53	**15.36 – 14.82**

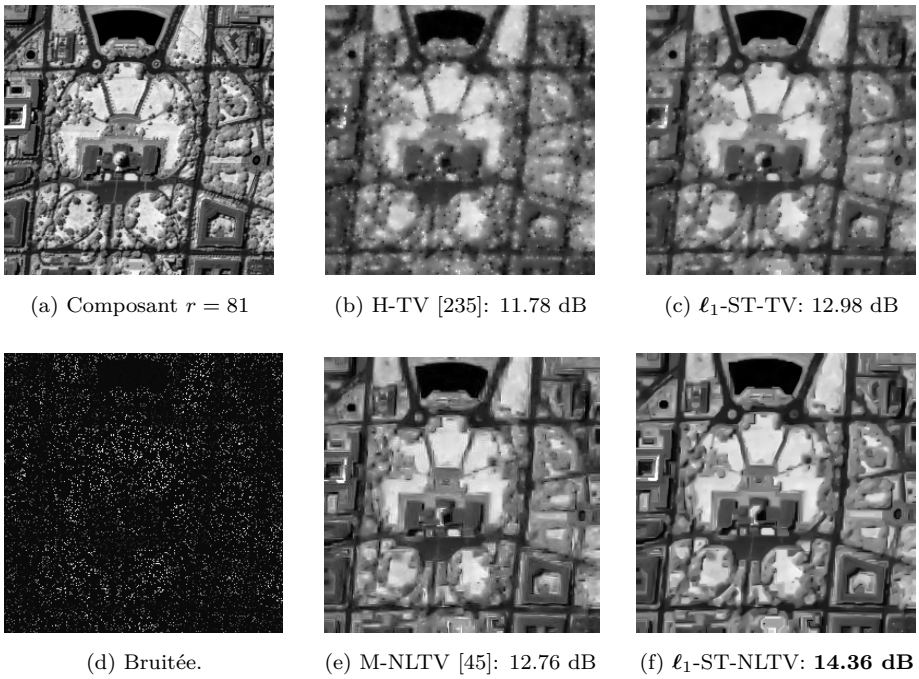

(a) Composant $r = 81$ (b) H-TV [235]: 11.78 dB (c) ℓ_1-ST-TV: 12.98 dB

(d) Bruitée. (e) M-NLTV [45]: 12.76 dB (f) ℓ_1-ST-NLTV: **14.36 dB**

Figure 8 *Comparaison d'une image hyperspectrale reconstruite avec H-TV, ℓ_1-ST-TV, M-NLTV et ℓ_1-ST-NLTV. Dégradation: scénario d'acquisition comprimée impliquant un bruit Gaussien additif avec moyenne nulle et écart type égal à* 5*, ainsi que* 90% *de décimation (*$N = 65536$*,* $R = 191$*,* $K = 6553$ *et* $S = 191$*).*

SVM PARCIMONIEUSE

Une machine à vecteurs de support (SVM) est une fonction qui prédit la classe $k \in \{1, \dots, K\}$ associée à une observation $u \in \mathbb{R}^N$, en utilisant des vecteurs de référence $\mathrm{x}^{(k)}$ tels que :

$$d(u) = \operatorname*{arg\,max}_{k \in \{1,\dots,K\}} \ \varphi(u)^\top \mathrm{x}^{(k)},$$

où $\varphi \colon \mathbb{R}^N \mapsto \mathbb{R}^{M+1}$. Le vecteur $\mathrm{x} = [\mathrm{x}^{(1)^\top} \ \dots \ \mathrm{x}^{(K)^\top}]^\top$ est estimé sur un ensemble de données $S = \Big\{(u_\ell, z_\ell) \in \mathbb{R}^N \times \{1, \dots, K\} \mid \ell = \{1, \dots, L\}\Big\}$ en résolvant un problème d'optimisation convexe :

$$\underset{\mathrm{x} \in \mathbb{R}^{(M+1)K}}{\text{minimize}} \ \sum_{k=1}^{K} \|x^{(k)}\|_2^2 + \lambda \sum_{\ell=1}^{L} \max\Big\{0, 1 + \max_{k \neq z_\ell} \ \varphi(u_\ell)^\top (\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)})\Big\}.$$

Nous proposons de remplacer la norme quadratique par une fonction convexe quelconque $f \in \Gamma_0(\mathbb{R}^{(M+1)K})$, ainsi que reformuler la fonction de perte ci-dessus en introduisant, pour chaque $\ell \in \{1, \dots, L\}$, l'opérateur linéaire $T_\ell \colon \mathbb{R}^{(M+1)K} \mapsto \mathbb{R}^K$ tel que

$$\Big(\forall \mathrm{x} \in \mathbb{R}^{(M+1)K}\Big) \qquad T_\ell \, \mathrm{x} = \begin{bmatrix} \varphi(u_\ell)^\top (\mathrm{x}^{(1)} - \mathrm{x}^{(z_\ell)}) \\ \cdots \\ \varphi(u_\ell)^\top (\mathrm{x}^{(K)} - \mathrm{x}^{(z_\ell)}) \end{bmatrix},$$

le vecteur $r_\ell = (r_\ell^{(k)})_{1 \leq k \leq K} \in \mathbb{R}^K$ tel que

$$(\forall k \in \{1, \dots, K\}) \qquad r_\ell^{(k)} = \begin{cases} 0, & \text{if } k = z_\ell, \\ \mu_\ell, & \text{otherwise,} \end{cases}$$

et la fonction $h_\ell \colon \mathbb{R}^K \mapsto \mathbb{R}$ telle que, pour chaque $\mathsf{y}^{(\ell)} = (y^{(\ell,k)})_{1 \leq k \leq K}$,

$$h_\ell(\mathsf{y}^{(\ell)}) = \max_{1 \leq k \leq K} y^{(\ell,k)} + r_\ell^{(k)},$$

de façon que $h_\ell(T_\ell \mathrm{x}) = \max\Big\{0, \mu_\ell + \max_{k \neq z_\ell} \ \varphi(u_\ell)^\top (\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)})\Big\}$. Notre objectif est donc de résoudre les problèmes suivants

$$\textit{approche regularisée:} \qquad \underset{\mathrm{x} \in \mathbb{R}^{(M+1)K}}{\text{minimize}} \quad f(\mathrm{x}) + \lambda \quad \sum_{\ell=1}^{L} h_\ell(T_\ell \, \mathrm{x}),$$

$$\textit{approche constrainte:} \qquad \underset{\mathrm{x} \in \mathbb{R}^{(M+1)K}}{\text{minimize}} \quad f(\mathrm{x}) \quad \text{s.t.} \ \sum_{\ell=1}^{L} h_\ell(T_\ell \, \mathrm{x}) \leq \eta,$$

où λ et η sont des constantes positives.

L'approche épigraphique permet de résoudre le problème sous contrainte au travers d'un vecteur $\zeta = \left(\zeta^{(\ell)}\right)_{1\leq\ell\leq L}$ tel que

$$\underset{(\mathrm{x},\zeta)\in\mathbb{R}^{(M+1)K}\times\mathbb{R}^{L}}{\text{minimize}} \quad f(\mathrm{x}) \qquad \text{s. t.} \qquad \begin{cases} (T\mathrm{x},\,\zeta) \in E, \\ \qquad \zeta \;\in V, \end{cases}$$

où

$$E = \Big\{(y,\zeta)\in\mathbb{R}^{LK}\times\mathbb{R}^{L} \;\Big|\; (\forall \ell\in\{1,\dots,L\}) \quad (y^{(\ell)},\zeta^{(\ell)})\in \operatorname{epi} h_\ell\Big\}, \tag{1}$$

$$V = \Big\{\zeta\in\mathbb{R}^{L} \;\Big|\; 1_L^\top\zeta\leq\eta\Big\}. \tag{2}$$

Les itérations de FBPD sont montrées dans l'algorithme 0.3. L'avantage de cette approche est que l'opérateur P_E équivaut à calculer des projections épigraphiques de la fonction max, pour laquelle nous avons dérivé une forme explicite.

Algorithm 0.3 *FBPD pour résoudre le problème sous contrainte.*

Initialization

$$\begin{cases} \text{choose } (\mathrm{x}^{[0]},\zeta^{[0]})\in\mathbb{R}^{(M+1)K}\times\mathbb{R}^{L} \\ \text{choose } (y^{[0]},\xi^{[0]})\in\mathbb{R}^{L(K-1)}\times\mathbb{R}^{L} \\ \text{set } \tau>0 \text{ and } \sigma>0 \text{ such that } \tau\sigma\max\{\|T\|^2,1\}\leq 1. \end{cases}$$

For $i = 0, 1, \dots$

$$\begin{cases} \mathrm{x}^{[i+1]} = \operatorname{prox}_{\tau f}\left(\mathrm{x}^{[i]} - \tau\, T^\top y^{[i]}\right) \\ \zeta^{[i+1]} = P_V\left(\zeta^{[i]} - \tau\,\xi^{[i]}\right) \\ \widehat{y}^{[i]} = y^{[i]} + \sigma T\left(2\mathrm{x}^{[i+1]} - \mathrm{x}^{[i]}\right) \\ \widehat{\xi}^{[i]} = \xi^{[i]} + \sigma\left(2\zeta^{[i+1]} - \zeta^{[i]}\right) \\ \left(\widetilde{y}^{[i]},\widetilde{\xi}^{[i]}\right) = P_E\left(\widehat{y}^{[i]}/\sigma,\widehat{\xi}^{[i]}/\sigma\right) \\ y^{[i+1]} = \widehat{y}^{[i]} - \sigma\widetilde{y}^{[i]} \\ \xi^{[i+1]} = \widehat{\xi}^{[i]} - \sigma\widetilde{\xi}^{[i]}. \end{cases}$$

Nous avons évalué les performances de notre approche sur trois bases de données, en considérant un scénario avec peu d'exemples. Les tests numériques montrent que l'utilisation d'une formulation exacte, au lieu d'une approximation, conduit à une meilleure précision de classification et à une méthode plus robuste par rapport au choix du paramètre de régularisation, tandis que l'algorithme proposé est souvent plus rapide que les solutions de l'état-de-l'art.

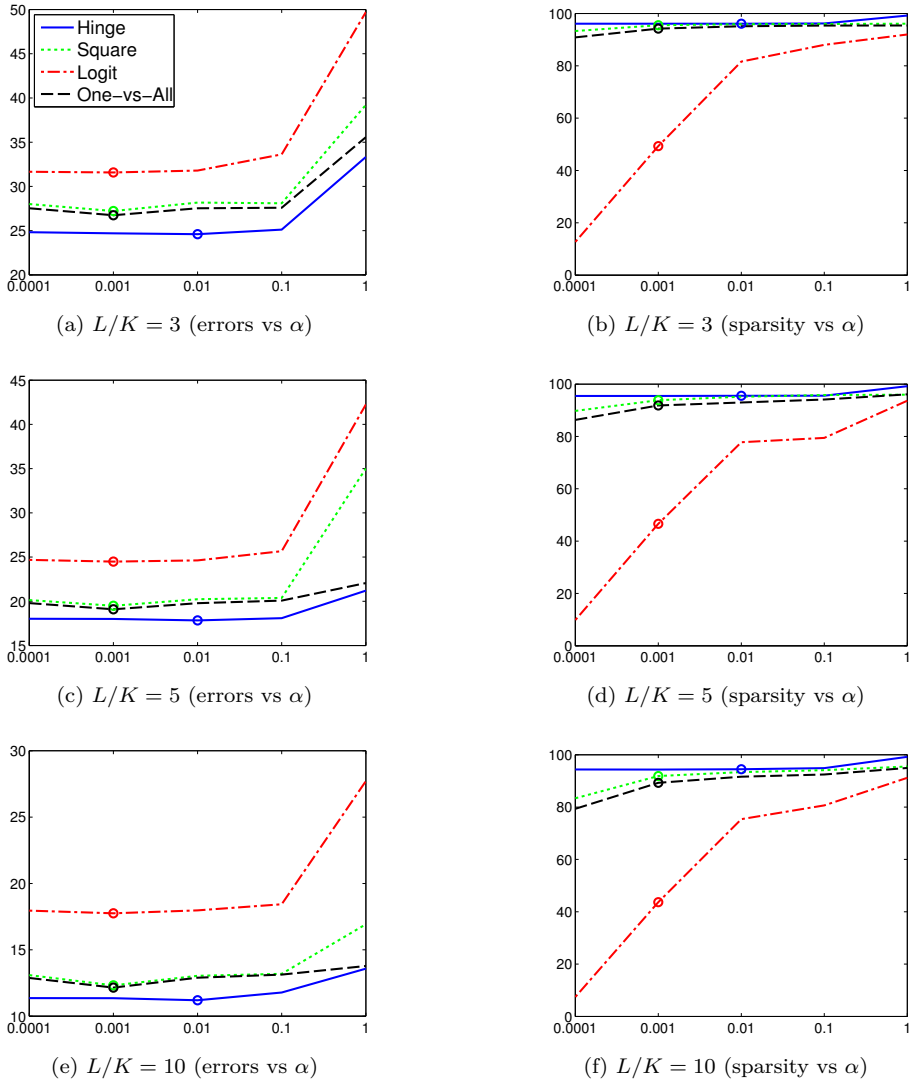

(a) $L/K = 3$ (errors vs α)

(b) $L/K = 3$ (sparsity vs α)

(c) $L/K = 5$ (errors vs α)

(d) $L/K = 5$ (sparsity vs α)

(e) $L/K = 10$ (errors vs α)

(f) $L/K = 10$ (sparsity vs α)

Figure 9 *Résultats sur le database MNIST pour $L \in \{3K, 5K, 10K\}$. Colonne de gauche: erreurs de classification en fonction de la régularisation. Colonne de droite: pourcentage de coefficients nuls.*

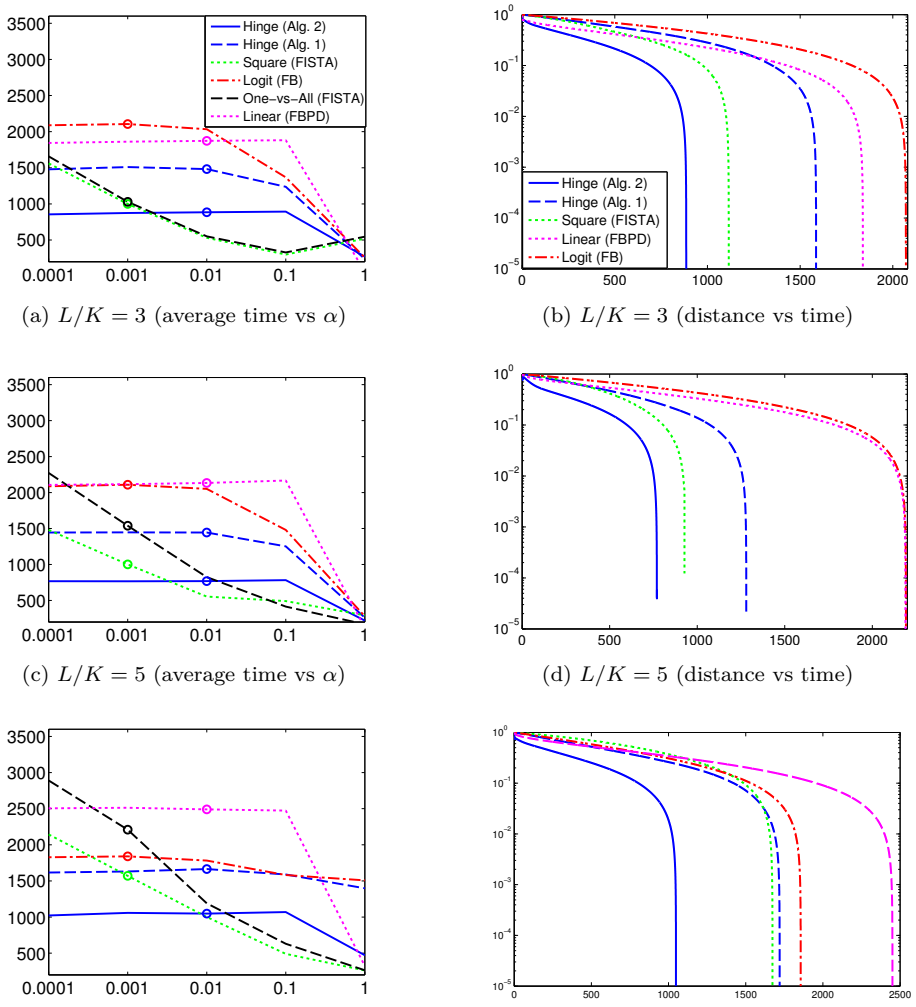

(a) $L/K = 3$ (average time vs α)

(b) $L/K = 3$ (distance vs time)

(c) $L/K = 5$ (average time vs α)

(d) $L/K = 5$ (distance vs time)

(e) $L/K = 10$ (average time vs α)

(f) $L/K = 10$ (distance vs time)

Figure 10 *Résultats sur le database MNIST pour $L \in \{3K, 5K, 10K\}$. Colonne de gauche: temps d'exécution en fonction de la régularisation. Colonne de droite: distance à la solution $\mathrm{x}^{[\infty]}$ en fonction du temps.*

Détection d'images falsifiées

Une image $y \in \mathbb{R}^N$ capturée par une camera digitale peut être décrite par un modèle de bruit additif-multiplicatif tel que

$$y = x \circ k + x + \theta,$$

où $\circ$ dénote le produit composant par composant, $x \in \mathbb{R}^N$ est l'image idéale sans bruit, $k \in \mathbb{R}^N$ est le PRNU de la camera, et $\theta \in \mathbb{R}^N$ est la réalisation d'un bruit additif. Ci-dessus, le terme k est le seul signal d'intérêt dans nos analyses, tandis que les autres termes jouent le rôle de bruits. Cela inclut l'image idéale x, qui est en fait estimée et soustraite de l'image originale, afin d'obtenir un résidu de bruit $r = y - \widehat{x}$ plus facile à traiter, où $\widehat{x} = f(y)$ est estimé par un filtre de debruitage f. For convenance, le résidu est réécrit de façon telle que k multiplie l'image observée y plutôt que l'image inconnue x :

$$r = y - \widehat{x} = y \circ k + \Big(x - y\Big) \circ k + \Big(x - \widehat{x}\Big) + \theta = y \circ k + n.$$

La procédure proposée en [41] pour vérifier l'intégrité d'une image est composée de trois étapes :

(i). **Estimation du PRNU de la camera** à partir d'un large nombre d'images non-falsifiées provenant de la même camera :

$$(\forall \ell \in \{1, \dots, N\}) \qquad \widehat{k}^{(\ell)} = \frac{\sum_{m=1}^{M} y_m^{(\ell)}\, r_m^{(\ell)}}{\sum_{m=1}^{M} \Big(y_m^{(\ell)}\Big)^2}.$$

(ii). **Calcul de la corrélation** entre le PRNU et le résidu de l'image

$$\rho^{(\ell)} = \operatorname{corr}\left(\Big[r^{(n)}\Big]_{n \in \mathcal{W}_\ell}, \Big[y^{(n)} k^{(n)}\Big]_{n \in \mathcal{W}_\ell}\right),$$

où $\mathcal{W}_\ell$ est le support du bloc 128×128 centré en ℓ.

(iii). **Test de falsification**, qui consiste à comparer la carte de corrélation calculée précédemment avec des seuils γ_1 et γ_2, afin d'obtenir une image binaire $\widehat{u}$ indiquant les zones falsifiées

$$\widehat{u}^{(\ell)} = \begin{cases} 0, & \text{si } \rho^{(\ell)} < \gamma_1 \text{ and } \widehat{\rho}^{(\ell)} > \gamma_2, \\ 1, & \text{sinon.} \end{cases}$$

Le résultat $\widehat{u}$ est post-traité avec les outils morphologiques.

Nous proposons de remplacer la troisième étape par la résolution d'un problème d'optimisation qui emploie une régularisation fondée sur les champs de Markov afin de prendre en compte les dépendances spatiales de la carte indiquant les zones falsifiées, conduisant à

$$\underset{u\in\{0,1\}^N}{\text{minimize}} \quad \sum_{\ell=1}^{N} c^{(\ell)}\, u^{(\ell)} + \lambda \sum_{\ell=1}^{N} \sum_{n\in\mathcal{N}_\ell} |u^{(\ell)} - u^{(n)}|,$$

où

$$c^{(\ell)} = \frac{\left(\rho^{(\ell)} - \widehat{\rho}^{(\ell)}\right)^2}{2\sigma_1^2} - \frac{\left(\rho^{(\ell)}\right)^2}{2\sigma_0^2} + \log \frac{p_0\,\sigma_1}{p_1\,\sigma_0}.$$

Le problème ci-dessus peut être résolu de manière exacte par une relaxation convexe, qui consiste à replacer le domaine discrète $\{0,1\}^N$ par l'hypercube unitaire $[0,1]^N$. De plus, la formulation convexe peut être réécrite sous forme contrainte

$$\underset{u\in[0,1]^N}{\text{minimize}} \quad \sum_{\ell=1}^{N} c^{(\ell)}\, u^{(\ell)} \quad \text{s. t.} \quad \sum_{\ell=1}^{N} \sum_{n\in\mathcal{N}_\ell} |u^{(\ell)} - u^{(n)}| \leq \eta,$$

ce qui peut simplifier le choix du paramètre η, car il est en lien direct avec la taille de la falsification. Le problème sous contrainte peut être résolu efficacement par l'éclatement épigraphique, à savoir

$$\underset{(u,\zeta)\in[0,1]^N\times V}{\text{minimize}} \quad c^\top u \quad \text{s. t.} \quad (Fx,\zeta)\in E,$$

où $c = (c^{(\ell)})_{1\leq\ell\leq N}$, F dénote l'opérateur différence, et

$$E = \Big\{(y,\zeta)\in\mathbb{R}^{4N}\times\mathbb{R}^{4N} \;\Big|\; (\forall \ell\in\{1,\dots,4N\}) \quad (y^{(\ell)},\zeta^{(\ell)})\in \operatorname{epi}|\cdot|\Big\}$$
$$V = \Big\{\zeta\in\mathbb{R}^{4N} \;\Big|\; 1_{4N}^\top\,\zeta\leq\eta\Big\}.$$

Les iterations de la méthode M+LFBF associée à ce problème sont montrées dans l'algorithme 0.4, où l'opérateur P_E équivaut à calculer des projections épigraphiques de la valeur absolue, pour laquelle nous avons dérivé une forme explicite.

Nous avons évalué les performance sur quatre cameras (Canon EOS 450D, Canon IXUS 95IS, Nikon D200, and Nikon Coolpix S5100) comprenant 600 images chacune. Les tests numériques montrent que notre approche abouti à une amélioration significative par rapport aux méthodes de l'état-de-l'art. Également, une version modifiée de l'algorithme proposé a permis à l'équipe GRIP de gagner le premier « IEEE IFS-TC Image Forensics Challenge ».

Algorithm 0.4 *Itérations de M+LFBF [66].*

INITIALIZATION

choose $\left(x^{[0]}, \zeta^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^{4N}$

choose $\left(y^{[0]}, \xi^{[0]}\right) \in \mathbb{R}^{4N} \times \mathbb{R}^{4N}$

set $\gamma \in]0, 1/4[$

FOR $i = 0, 1, \ldots$

$$
\begin{aligned}
p^{[i]} &= P_{[0,1]^N}\left(x^{[i]} - \gamma\left(c + F^\top y^{[i]}\right)\right) \\
\rho^{[i]} &= P_V\left(\zeta^{[i]} - \gamma\,\xi^{[i]}\right) \\
\widehat{y}^{[i]} &= y^{[i]} + \gamma F x^{[i]} \\
\widehat{\xi}^{[i]} &= \xi^{[i]} + \gamma\,\zeta^{[i]} \\
(v^{[i]}, \nu^{[i]}) &= (\widehat{y}^{[i]}, \widehat{\xi}^{[i]}) - \gamma P_E\left(\widehat{y}^{[i]}/\gamma, \widehat{\xi}^{[i]}/\gamma\right) \\
y^{[i+1]} &= v^{[i]} + \gamma F\left(p^{[i]} - x^{[i]}\right) \\
\xi^{[i+1]} &= \nu^{[i]} + \gamma\left(\rho^{[i]} - \zeta^{[i]}\right) \\
x^{[i+1]} &= p^{[i]} - \gamma F^\top (v^{[i]} - y^{[i]}) \\
\zeta^{[i+1]} &= \rho^{[i]} - \gamma\left(\nu^{[i]} - \xi^{[i]}\right)
\end{aligned}
$$

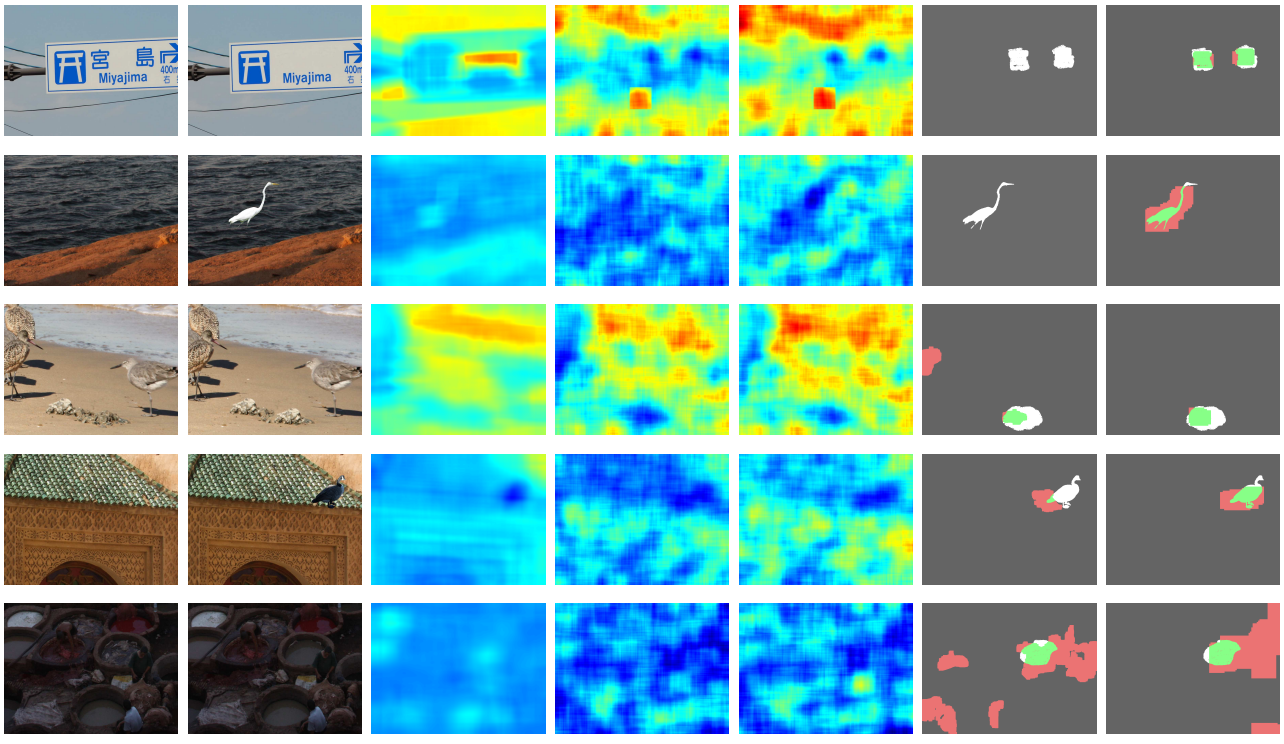

Originale | Falsifiée | Prediction $\widehat{\rho}$ | Corrélation ρ (avec Mihcak) | Corrélation ρ (avec BM3D) | Etat-de-l'art | Proposé

Figure 11 *Exemples d'images falsifiées et leur détection.*

Conclusion et perspectives

Nous avons proposé une technique d'éclatement épigraphique qui permet de réduire la complexité des algorithmes d'optimisation en présence d'une contrainte impliquant la somme de fonctions convexes telles que la valeur absolue, la fonction distance, la norme Euclidienne, la norme infinie et la fonction max. Nous avons démontré l'efficience de cette approche sur des problèmes ayant une contrainte sur la norme mixte, la divergence de Kullback-Leibler ou la régression logistique. Nous avons également apporté des contributions dans le contexte de la reconstruction d'images multi-composants, de l'apprentissage d'une SVM parcimonieuse et de la détection d'images falsifiées.

Le travail de cette thèse a ouvert des perspectives intéressantes. Une de celles-ci consiste à combiner les approches d'optimisation stochastique avec l'éclatement épigraphique. Cela est de grand intérêt dans toutes les applications impliquant l'élaboration d'un volume important de données, comme par exemple l'apprentissage d'un classificateur parcimonieux. A l'état actuel, les algorithmes existants (dont ceux que j'ai proposés) sont limités par le nombre d'échantillons sur lesquels l'apprentissage peut être effectué. L'approche stochastique permettrait de surmonter cette problématique et de rendre nos algorithmes d'apprentissage parcimonieux encore plus compétitifs par rapport aux techniques classiques d'apprentissage, avec des retombées applicatives très intéressantes.

Une deuxième perspective serait d'étendre la régularisation basée sur le tenseur de structure aux applications d'imagerie 3D. Un exemple concerne les images volumétriques, qui sont largement utilisées dans le contexte biomédical. Le tenseur de structure permettrait d'adapter la restauration d'images aux besoins spécifiques de cette communauté, par exemple en assurant de préserver les textures complexes et les structure fines, telles que les vaisseaux sanguins. Bien évidemment, cette approche peut s'étendre à d'autres domaines, comme par exemple l'imagerie SAR multi-temporelle.

Enfin, une troisième perspective concerne l'allocation optimale des ressources dans un codeur vidéo hybride. Dans un travail préliminaire [98], nous avons traité un problème d'optimisation convexe qui vise à minimiser la distorsion moyenne, tout en respectant une contrainte de budget sur les ressources allouées aux différentes trames d'un vidéo. Une extension de ce travail consiste à formuler le problème sous forme contrainte, afin d'avoir un contrôle direct sur la distorsion du vidéo codifié. L'approche épigraphique se présente alors comme une solution naturelle, et nous nous attendons des performances cohérentes aux résultats présentés dans cette thèse.

Part I

Methodology

Life is short and information endless: nobody has time for everything. In practice we are generally forced to choose between an unduly brief exposition and no exposition at all. Abbreviation is a necessary evil and the abbreviator's business is to make the best of a job which, although intrinsically bad, is still better than nothing.

Aldous Huxley

Chapter 1
Convex optimization

The main focus of this chapter is to review the basics of non-smooth convex optimization. We begin with a brief overview of inverse problems addressed by variational approaches. Then, we present a panel of convex optimization methods based on proximal tools. Finally, we discuss different techniques for computing the projection onto a convex set.

1.1 Inverse problems

In many branches of science, one is interested in converting some measurements into information about an entity that cannot be directly observed. This task can be considered as the problem of inverting the forward relation between the model parameters and the observed data. A popular example is the task of recovering a signal $\overline{x} \in \mathbb{R}^N$ from an observation $z \in \mathbb{R}^K$ given by the linear model

$$z = A\,\overline{x} + b, \tag{1.1}$$

where $A \in \mathbb{R}^{K \times N}$ is known, and $b \in \mathbb{R}^K$ is a realization of a zero-mean noise. Unfortunately, even in the absence of noise ($b = 0$), there is no guarantee that the solution to this problem is unique or even exists, as the above system may be underdetermined ($\mathrm{rank}(A) < N$) or overdetermined ($\mathrm{rank}(A) = N$ and $K > N$).

In general, inverse problems can be hardly solved in an exact manner, due to uncertainties in the observed data, and issues related to existence, uniqueness, and stability of the solution. For such ill-posed problems, one actually aims at finding inexact solutions which are optimal in a statistical or application-specific sense. As a result, the standard approach to deal with inverse problems involves the use of mathematical optimization, like in the celebrated *least-squares method*.

1.1.1 Least-squares method

The method of least squares is a standard approach to find the approximate solution to an overdetermined system. The objective is to recover the signal that minimizes the quadratic distance to the observation, leading to the problem

$$\underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \|Ax - z\|_2^2. \tag{1.2}$$

The above problem is solved for $\widehat{x} = (A^\top A)^{-1} A^\top z$. However, when the matrix A is ill conditioned, this solution is highly unstable and sensitive to small changes in the observed data. In such a case, it is necessary to incorporate additional information in order to prefer a solution with desirable properties. A possible approach consists of adding a quadratic regularization [209], yielding

$$\underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \|Ax - z\|_2^2 + \lambda \|Fx\|_2^2, \tag{1.3}$$

where $\lambda > 0$ and $F \in \mathbb{R}^{M \times N}$. can be the identity matrix, a gradient operator, etc. So doing, the solution changes to $\widehat{x} = (A^\top A + \lambda F^\top F)^{-1} A^\top z$, which is well defined

in the case when the matrix $[A^\top \; F^\top]^\top$ has full column rank.

An alternative regularization for the least-squares method amounts to introducing a constraint that bounds the $\boldsymbol{\ell}_1$-norm of the solution [208, 217], yielding

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad \|Ax - z\|_2^2 \quad \text{s.t.} \quad \|Fx\|_1 \leq \eta, \tag{1.4}$$

with $\eta > 0$. The $\boldsymbol{\ell}_1$-norm ensures that the solutions will have a number of transformed coefficients exactly equal to zero, depending on the value of η. Such a sparsity-inducing behavior is exploited, for example, in the context of image recovery [2, 8, 40, 80, 81, 195, 207]. However, unlike the quadratic regularization, the above problem cannot be solved in closed form. This is typical of inverse problems, which generally need to be solved through optimization methods.

1.1.2 Single-objective optimization

Optimization aims at selecting the best element from a set of available alternatives ranked by a given criterion. This problem can be formally stated as

$$\underset{x\in D}{\text{minimize}} \quad f(x), \tag{1.5}$$

The function f can take the value $+\infty$ for discarding some "forbidden parts" of D.

where D denotes the set of feasible solutions, and $f\colon D \mapsto \mathbb{R} \cup \{+\infty\}$ denotes the objective function. Optimization problems can be divided into two categories depending on whether the feasible set is continuous ($D \subset \mathbb{R}^N$) or discrete ($D \subset \mathbb{Z}^N$). In the former case, the feasible solutions may be signals or images, whereas in the latter case they may be permutations or graphs.

There is no general formula for the solution to optimization problems. Hence, they are solved by generating a sequence of points converging towards an optimal solution. In this regard, variational techniques are used to deal with continuous problems, as well as search techniques and cutting-plane methods are used to tackle discrete problems. Among the huge array of optimization methods, many of them have been specifically developed for solving convex problems. These algorithms will be throughly discussed in Section 1.2. Given the efficiency and the reliability of convex optimization methods, they are often employed to approach nonconvex problems through the use of convex relaxations. This applies for example to discrete labeling problems [36, 203], linear integer problems [135], and continuous problems involving certain nonconvex functions [34, 173, 184].

1.1.3 Multi-objective optimization

Inverse problems typically involve multiple objectives to be optimized, such as data fidelity and regularization terms, as well as criteria enforcing additional pieces of prior information. This leads to a multi-objective optimization problem:

$$\underset{x\in D}{\text{minimize}} \quad \Big\{f_1(x), \ldots, f_R(x)\Big\}, \tag{1.6}$$

with $f_r\colon D \mapsto \mathbb{R} \cup \{+\infty\}$ for every $r \in \{1, \ldots, R\}$. In multi-objective optimization, there hardly exists a solution that minimizes all objective functions simultaneously, hence attention is paid to *Pareto optimal solutions* [169], namely the solutions for which it is impossible to decrease the value of any objective function without increasing the value of at least another objective function.

A possible approach to deal with a multi-objective optimization problem consists of converting it into a single-objective problem such that its optimal solutions are some of the Pareto optimal solutions to the former. In this regard, the most well-known reformulations of Problem (1.6) are the following [169]:

- **regularized formulation**

$$\underset{x\in D}{\text{minimize}} \quad \lambda_1 f_1(x) + \cdots + \lambda_R f_R(x), \tag{1.7}$$

  with $\lambda_r > 0$ for every $r \in \{1, \ldots, R\}$;

- **constrained formulation**

$$\underset{x\in D}{\text{minimize}} \quad f_j(x) \quad \text{s.t.} \quad (\forall r \neq j) \quad f_r(x) \leq \eta_r, \tag{1.8}$$

  with $\eta_r \in \mathbb{R}$ such that $\left\{x \;\middle|\; f_r(x) \leq \eta_r\right\} \neq \emptyset$ for every $r \in \{1, \ldots, R\} \setminus \{j\}$.

Under technical assumptions [138, 181], the above formulations are equivalent for some specific values of $(\lambda_r)_{1\leq r\leq R}$ and $(\eta_r)_{1\leq r\leq R}$. However, different Pareto optimal solutions are produced by varying these parameters, leading to results with a different quality. There is no general rule that relates the parameters $(\lambda_r)_{1\leq r\leq R}$ to the quality of solutions obtained with the regularized formulation, although it is possible to optimally select them in some circumstances [82, 87, 93, 159, 179]. In this regard, the constrained formulation may be more practical, in the sense that the bounds $(\eta_r)_{1\leq r\leq R}$ may be easier to select, as they are related to some knowledge on the degradation process [2, 67, 207] or some physical properties of the target solution [61, 217]. This is the main motivation in considering constrained formulations of optimization problems.

1.2 Proximal algorithms

The first methods for finding a solution to an inverse problem were restricted to the use of a differentiable function [209]. In this context, gradient-based algorithms (e.g., nonlinear conjugate gradient or quasi-Newton methods) are popular [54]. However, in order to model properties such as sparsity or dynamic ranges, one needs to incorporate non-smooth functions or hard constraints in the optimization problem. In this case, a possible approach consists of resorting to smart approximations in order to smooth the involved non-differentiable functions [11, 14, 115, 123]. If one wants to directly address the original non-smooth problem, projection methods [35, 232], block-coordinate descent methods [212], or interior-point methods [225] may be applied, but they often impose restrictive assumptions on the problem to be solved. To overcome these issues, a new class of iterative methods referred to as *proximal algorithms* has recently emerged to efficiently solve a large panel of convex optimization problems.

1.2.1 Proximity operator

The key tool in proximal methods is the *proximity operator* [174] of a proper lower semicontinuous convex function $f \in \Gamma_0(\mathbb{R}^N)$, defined as

$\Gamma_0(\mathbb{R}^N)$ *is the class of proper lower semicontinuous convex functions from* $\mathbb{R}^N$ *to* $\mathbb{R}\cup\{+\infty\}$.

$$(\forall x \in \mathbb{R}^N) \qquad \operatorname{prox}_f(x) = \underset{u\in\mathbb{R}^N}{\arg\min} \; \frac{1}{2}\|u - x\|_2^2 + f(u). \tag{1.9}$$

The subgradient of $\varphi \in \Gamma_0(\mathbb{R}^N)$ at $p \in \mathbb{R}^N$ is the set defined as

$$\partial\varphi(p) = \left\{ t \in \mathbb{R}^N \mid (\forall u \in \mathbb{R}^N) \; (u-p)^\top t + \varphi(p) \leq \varphi(u) \right\}.$$

If φ is differentiable at y, then

$$\partial\varphi(y) = \{\nabla\varphi(y)\}.$$

The proximity operator can be interpreted as a sort of implicit subgradient step for the function f, as it is uniquely defined through the following inclusion

$$p = \operatorname{prox}_f(x) \quad \Leftrightarrow \quad p \in x - \partial f(p). \tag{1.10}$$

Proximity operators enjoy many interesting properties [12, 40, 65, 178].

The great advantage of proximal methods is that they provide a unifying framework which allows one to deal with non-smooth functions as well as hard constraints. Indeed, the constraint associated with a nonempty closed convex subset $C \subset \mathbb{R}^N$ can be enforced through the indicator function of C, defined as

$$(\forall x \in \mathbb{R}^N) \qquad \iota_C(x) = \begin{cases} 0, & \text{if } x \in C, \\ +\infty, & \text{otherwise.} \end{cases} \tag{1.11}$$

This function belongs to $\Gamma_0(\mathbb{R}^N)$ and the associated proximity operator coincides to the *orthogonal projection* onto C, in the sense that

$$(\forall x \in \mathbb{R}^N) \qquad \operatorname{prox}_{\iota_C}(x) = P_C(x) = \operatorname*{arg\,min}_{u \in C} \; \|u - x\|_2^2. \tag{1.12}$$

Proximal algorithms can be classified in primal [1, 65, 180] and primal-dual [29, 37, 66, 69, 135, 215] methods. The main difference is that the latter ones do not need to invert the linear operators involved in the optimization problem, even though the former ones may converge faster when such an inversion can be efficiently computed. Interestingly, a number of primal-dual methods can be derived from the *forward-backward* splitting [59].

1.2.2 Forward-backward splitting

The gradient of $\varphi \in \Gamma_0(\mathbb{R}^N)$ is β-Lipschitz if there exists $\beta > 0$ such that: $\forall (x,u) \in \mathbb{R}^N \times \mathbb{R}^N$

$$\|\nabla\varphi(x) - \nabla\varphi(u)\|_2 \leq \beta \|x - u\|_2.$$

Forward-backward splitting is a standard approach to optimize the sum of a non-smooth function $f \in \Gamma_0(\mathbb{R}^N)$ and a differentiable function $g \in \Gamma_0(\mathbb{R}^N)$ with a β-Lipschitz continuous gradient for some $\beta \in]0, +\infty[$, namely

$$\operatorname*{minimize}_{x \in \mathbb{R}^N} \quad f(x) + g(x). \tag{1.13}$$

This method aims at finding a point $\overline{x} \in \mathbb{R}^N$ satisfying the fixed-point condition

$$0 \in \partial f(\overline{x}) + \nabla g(\overline{x}) \quad \Leftrightarrow \quad \overline{x} \in \overline{x} - \nabla g(\overline{x}) - \partial f(\overline{x}) \quad \Leftrightarrow \quad \overline{x} = \operatorname{prox}_f\Big(\overline{x} - \nabla g(\overline{x})\Big). \tag{1.14}$$

As a matter of fact, under appropriate technical assumptions, it can be shown [68] that, for every $x^{[0]} \in \mathbb{R}^N$, the sequence $(x^{[i]})_{i \in \mathbb{N}}$ generated by

$$x^{[i+1]} = \operatorname{prox}_{\gamma_i f}\Big(x^{[i]} - \gamma_i \nabla g(x^{[i]})\Big) \tag{1.15}$$

converges to a solution to Problem (1.13) for any $\gamma_i \in]0, 2/\beta[$.

The inf-convolution between φ and ψ in $\Gamma_0(\mathbb{R}^N)$ is defined as

$$\varphi \,\square\, \psi = \inf_{u \in \mathbb{R}^N} \; \varphi(u) + \psi(\cdot - u).$$

It is differentiable for $\psi = \|\cdot\|_2^2$

$$\nabla(\varphi \,\square\, \tfrac{1}{2\gamma}\|\cdot\|_2^2) = (\cdot - \operatorname{prox}_{\gamma\varphi})/\gamma.$$

The forward-backward splitting generalizes various well-known algorithms, which can be recovered from (1.15) through specific choices of f and g. This includes the standard gradient descent when $f = 0$, the proximal-point algorithm when $g = 0$, the projected gradient method when f is the indicator function of a closed convex subset of $\mathbb{R}^N$, and the backward-backward splitting when g is the inf-convolution between a function in $\Gamma_0(\mathbb{R}^N)$ and the quadratic function.

1.2.3 Fenchel-Rockafellar duality

A strong practical limitation of forward-backward splitting arises when the non-smooth term f denotes a convex function composed with a linear operator. Indeed, the proximity operator of such a function is difficult to compute, unless the linear operator satisfies specific properties [12, Propos. 23.23], [186, Propos. 3.4]. One way to circumvent this issue is provided by the *Fenchel-Rockafellar duality*, which can be applied to Problem (1.13) with an additional function $h \in \Gamma_0(\mathbb{R}^M)$ composed with a linear operator $F \in \mathbb{R}^{M\times N}$, and some $w \in \mathbb{R}^N$, namely

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad f(x) + g(x) + h(Fx) - w^\top x. \tag{1.16}$$

Let $F\in\mathbb{R}^{M\times N}$ and $\varphi\in\Gamma_0(\mathbb{R}^M)$. If $FF^\top = \nu I$ with $\nu > 0$, then

$$\operatorname{prox}_{\varphi\circ F}(x) = x + \tfrac{1}{\nu}F^\top\big(\operatorname{prox}_{\nu\varphi}(Fx) - Fx\big).$$

If $F^\top = F^{-1}$, then

$$\operatorname{prox}_{\varphi\circ F}(x) = F^\top \operatorname{prox}_{\varphi}(Fx).$$

The principle is that the above *primal problem* can be associated with a *dual problem* involving the conjugate functions of $f+g$ and h, yielding [12, Definition 15.19]

$$\underset{y\in\mathbb{R}^M}{\text{minimize}} \quad (f^* \,\square\, g^*)(w - F^\top y) + h^*(y). \tag{1.17}$$

Let φ and ψ belong to $\Gamma_0(\mathbb{R}^N)$. The conjugate function of φ is

$$\varphi^* = \sup_{x\in\mathbb{R}^N} x^\top(\cdot) - \varphi(x).$$

In addition, $(\varphi+\psi)^* = \varphi^* \,\square\, \psi^*$, and $u \in \partial\varphi(x) \;\Leftrightarrow\; x \in \partial\varphi^*(u) \;\Leftrightarrow\; \varphi(x) + \varphi^*(u) = x^\top u$.

Under some technical assumptions, the primal and dual problems are equivalent, in the sense that their objective functions achieve the same (up to a sign) optimum value [12, Theorem 15.23]. This fact allows one to deduce that any pair $(\overline{x}, \overline{y})$ of solutions to these problems is such that [12, Theorem 19.1]

$$\begin{cases} \overline{x} \in \partial(f^* \,\square\, g^*)(w - F^\top \overline{y}), \\ F\overline{x} \in \partial h^*(\overline{y}). \end{cases} \tag{1.18}$$

In addition, if there exists a solution $\overline{y}$ to the dual problem such that $f^* \,\square\, g^*$ is differentiable at $w - F^\top\overline{y}$, then the primal problem admits either no solution or a unique solution such that $\overline{x} = \nabla(f^* \,\square\, g^*)(w - F^\top\overline{y})$ [12, Proposition 19.3].

An important application of Fenchel-Rockafellar duality concerns the strongly convex problems. Indeed, when $g = \frac{1}{2}\|\cdot\|_2^2$, it is easier to solve the dual problem and to recover from it the (unique) solution to the primal problem, as

(i). the term $f^* \,\square\, \frac{1}{2}\|\cdot\|_2^2$ is differentiable with gradient $\nabla(f^* \,\square\, \frac{1}{2}\|\cdot\|_2^2) = \operatorname{prox}_f$,

(ii). there exists at least one solution $\overline{y}$ to the dual problem, and the solution $\overline{x}$ to the primal problem is such that $\overline{x} = \operatorname{prox}_f(w - F^\top\overline{y})$,

(iii). in the dual problem, the linear operator $F^\top$ appears in the smooth term $f^* \,\square\, \frac{1}{2}\|\cdot\|_2^2$ rather than the nonsmooth term h^*.

These arguments can be exploited to solve the dual problem via the forward-backward splitting. Under some technical assumptions, it can be shown [60] that, for every $y^{[0]} \in \mathbb{R}^M$, the sequence $(x^{[i]})_{i\in\mathbb{N}}$ generated by

$$\left\lfloor \begin{aligned} x^{[i]} &= \operatorname{prox}_f(w - F^\top y^{[i]}) \\ y^{[i+1]} &= \operatorname{prox}_{\gamma_i h^*}(y^{[i]} + \gamma_i F x^{[i]}) \end{aligned} \right. \tag{1.19}$$

converges to the solution to (1.16) with $g = \frac{1}{2}\|\cdot\|_2^2$, for any $\gamma_i \in \left]0, 2/\|F\|^2\right[$.

The proximity operator of φ and φ^* in $\Gamma_0(\mathbb{R}^N)$ are such that: for all $u \in \mathbb{R}^N$ and $\gamma > 0$

$$\operatorname{prox}_{\gamma\varphi^*}(u) = u - \gamma \operatorname{prox}_{\frac{\varphi}{\gamma}}\left(\frac{u}{\gamma}\right).$$

The above algorithm has the ability to decompose the optimization process into elementary steps such as prox_f, $\operatorname{prox}_{h^*}$, F, and $F^\top$. This makes it possible to reduce the complexity related to Problem (1.16) in the case when $g = \frac{1}{2}\|\cdot\|_2^2$. Fortunately, the same kind of decomposition can be obtained for general instances of Problem (1.16) through the use of *primal-dual* approaches.

1.2.4 Primal-dual approach

A more general way to employ Fenchel-Rockafellar duality consists of solving the primal-dual problem associated with (1.16), which can be derived from the latter by rewriting h in terms of its conjugate, yielding

A function $\varphi \in \Gamma_0(\mathbb{R}^N)$ can be expressed in terms of φ^* as

$$\varphi = \sup_{u \in \mathbb{R}^N} u^\top(\cdot) - \varphi^*(u).$$

$$\underset{y \in \operatorname{dom} h^*}{\text{maximize}} \quad \underset{x \in \mathbb{R}^N}{\text{minimize}} \quad f(x) + g(x) + y^\top F x - h^*(y) - w^\top x, \tag{1.20}$$

where $\operatorname{dom} h^* = \big\{ y \in \mathbb{R}^M \;\big|\; h^*(y) < +\infty \big\}$. So doing, one actually aims at solving simultaneously both the primal and dual problems in (1.16)-(1.17), leading to

$$\begin{cases} \overline{x} \in \underset{x \in \mathbb{R}^N}{\arg\min}\; f(x) + g(x) - (w - F^\top \overline{y})^\top x, \\ \overline{y} \in \underset{y \in \mathbb{R}^M}{\arg\min}\; h^*(y) - y^\top F \overline{x}. \end{cases} \tag{1.21}$$

The goal of primal-dual methods is to find a pair $(\overline{x}, \overline{y}) \in \mathbb{R}^N \times \mathbb{R}^M$ such that

$$\begin{cases} 0 \in \partial f(\overline{x}) + \nabla g(\overline{x}) - w + F^\top \overline{y} \\ 0 \in \partial h^*(\overline{y}) - F\overline{x}, \end{cases} \quad \Leftrightarrow \quad \begin{cases} \overline{x} = \operatorname{prox}_f \big(\overline{x} - \nabla g(\bar{x}) + w - F^\top \overline{y} \big) \\ \overline{y} = \operatorname{prox}_{h^*} \Big(\overline{y} + F\overline{x} \Big). \end{cases} \tag{1.22}$$

Under technical assumptions, the above condition guarantees that $\overline{x}$ is a solution to (1.16), $\overline{y}$ is a solution to (1.17), and $(\overline{x}, \overline{y})$ is a solution to (1.20) [12, Theorem 19.1, Proposition 19.18]. The advantage of (1.22) is that Problem (1.16) can be solved through the separate evaluation of operators prox_f, ∇g, and prox_h, which naturally leads to algorithms easier to implement. In addition, the Lipschitz-differentiable function is activated through its gradient, which is often simpler to compute than its proximity operator. Note also that Problem (1.16) can actually include an arbitrary number of non-smooth functions by setting

$$F = [F_1^\top \; \dots \; F_R^\top]^\top, \qquad y = [y_1^\top \; \dots \; y_R^\top]^\top, \qquad h(y) = h_1(y_1) + \cdots + h_R(y_R),$$

with $F_r \in \mathbb{R}^{M_r \times N}$, $y_r \in \mathbb{R}^{M_r}$, and $h_r \in \Gamma_0(\mathbb{R}^{M_r})$ for every $r \in \{1, \dots, R\}$.

Several primal-dual methods can be derived from a more general version of the forward-backward splitting [59], such as the one illustrated in Algorithm 1.1. There also exist primal-dual methods based on other approaches, like the one in Algorithm 1.2. The reader can refer to [135] for a survey on primal-dual methods.

Algorithm 1.1 *FBPD method [69, 215]*

Initialization

choose $\big(x^{[0]}, y^{[0]}\big) \in \mathbb{R}^N \times \mathbb{R}^M$

set $\tau > 0$ and $\sigma > 0$ such that $\tau \big(\beta/2 + \sigma \|F\|^2\big) < 1$

For $i = 0, 1, \dots$

$\widehat{x}^{[i]} = \nabla g(x^{[i]}) - w + F^\top y^{[i]}$

$x^{[i+1]} = \operatorname{prox}_{\tau f} \big(x^{[i]} - \tau\, \widehat{x}^{[i]} \big)$

$\widehat{y}^{[i]} = F\big(2x^{[i+1]} - x^{[i]}\big)$

$y^{[i+1]} = \operatorname{prox}_{\sigma h^*} \big(y^{[i]} + \sigma\, \widehat{y}^{[i]} \big)$

Algorithm 1.2 *M+LFBF method [66]*

Initialization

choose $\big(x^{[0]}, y^{[0]}\big) \in \mathbb{R}^N \times \mathbb{R}^M$

set $\gamma \in \left]0, (\beta + \|F\|)^{-1}\right[$

For $i = 0, 1, \dots$

$\widehat{x}^{[i]} = \nabla g(x^{[i]}) - w + F^\top y^{[i]}$

$p^{[i]} = \operatorname{prox}_{\gamma f} \big(x^{[i]} - \gamma\, \widehat{x}^{[i]} \big)$

$v^{[i]} = \operatorname{prox}_{\gamma h^*} \big(y^{[i]} + \gamma\, F x^{[i]} \big)$

$y^{[i+1]} = v^{[i]} + \gamma F \big(p^{[i]} - x^{[i]} \big)$

$\widetilde{x}^{[i]} = \nabla g(p^{[i]}) - w + F^\top v^{[i]}$

$x^{[i+1]} = p^{[i]} - \gamma \big(\widetilde{x}^{[i]} - \widehat{x}^{[i]} \big)$

1.3 Projection operators

Proximal methods deal with Problem (1.16) by iterating a sequence of steps in which the operators prox_f and prox_h are evaluated at each iteration. The efficient computation of these operators is thus essential for dealing with large-size convex optimization problems. While the proximity operator can be efficiently computed for many functions of practical interest [40, 65, 178], the same cannot be claimed when a hard constraint is involved, because the projection operator rarely admits a closed-form expression, except in the cases listed in Section 1.3.1. As a result, the resolution of constrained optimization problems such as (1.8) is often more challenging than solving the regularized counterparts in (1.7), due to the difficulty of computing the involved projections. Two popular approaches to circumvent this issue are discussed in Sections 1.3.2 and 1.3.3.

1.3.1 Projections in closed form

There exists a few number of convex sets for which an expression of the associated projection is available. The most popular and relevant examples are listed below.

Euclidean balls

Let $\varepsilon \in]0, +\infty[$, let $C \subset \mathbb{R}^N$ be a nonempty closed convex set, and let

$$B_\varepsilon = \big\{x \in \mathbb{R}^N \;\big|\; \|x\|_2 \leq \varepsilon\big\}. \tag{1.23}$$

The projection onto the convex set $B_\varepsilon + C$ reads [12, Proposition 28.10]

$$(\forall x \in \mathbb{R}^N) \qquad P_{B_\varepsilon + C}(x) = \begin{cases} x, & \text{if } \|x - P_C(x)\|_2 \leq \varepsilon, \\ P_C(x) + \varepsilon \, \dfrac{x - P_C(x)}{\|x - P_C(x)\|_2}, & \text{otherwise.} \end{cases} \tag{1.24}$$

Note that $B_\varepsilon + C$ reduces to the $\boldsymbol{\ell}_2$-ball centered in some $z \in \mathbb{R}^N$ when $C = \{z\}$.

Hyperslabs

Let $u \in \mathbb{R}^N \setminus \{0\}$, let $(\eta_1, \eta_2) \in \mathbb{R}^2$ such that $\eta_1 \leq \eta_2$, and let

$$C = \big\{x \in \mathbb{R}^N \;\big|\; \eta_1 \leq u^\top x \leq \eta_2\big\}. \tag{1.25}$$

The associated projection can be expressed as [12, Example 28.17]

$$(\forall x \in \mathbb{R}^N) \qquad P_C(x) = \begin{cases} x + \dfrac{\eta_1 - u^\top x}{\|u\|_2^2} u, & \text{if } u^\top x < \eta_1, \\ x + \dfrac{\eta_2 - u^\top x}{\|u\|_2^2} u, & \text{if } u^\top x > \eta_2, \\ x, & \text{otherwise.} \end{cases} \tag{1.26}$$

As a special case, C reduces to a closed half-space when $\eta_1 = -\infty$, an hyperplane when $\eta_1 = \eta_2$, and a closed interval when $N = 1$ and $u \equiv 1$.

Affine sets

Let $A \in \mathbb{R}^{K \times N}$. The *range* of A is the linear span of the columns of A, namely

$$\operatorname{ran}(A) = \big\{ z \in \mathbb{R}^K \;\big|\; (\exists\, x \in \mathbb{R}^N) \quad z = Ax \big\}. \tag{1.27}$$

The associated projection can be expressed as [12, Proposition 3.25]

$$(\forall z \in \mathbb{R}^K) \qquad P_{\operatorname{ran}(A)}(z) = AA^\dagger z, \tag{1.28}$$

where $A^\dagger$ denotes the pseudoinverse of A [12, Definition 3.26]. Note that $A^\dagger z$ yields a solution to the least-squares problem when A has full column rank, as

In the case when $\operatorname{rank}(A) = K$, $P_{\operatorname{ran}(A)}(z) = z$ *as* $\operatorname{ran}(A) = \mathbb{R}^K$.

$$A^\dagger = \begin{cases} (A^\top A)^{-1} A^\top, & \text{if } \operatorname{rank}(A) = N, \\ A^\top (AA^\top)^{-1}, & \text{if } \operatorname{rank}(A) = K. \end{cases} \tag{1.29}$$

Moreover, the *kernel* of A is the orthogonal complement of the range of $A^\top$, i.e.

For a vector subspace $V \subset \mathbb{R}^N$ $(\iota_V)^* = \iota_{V^\perp}$, *which yields: for every* $x \in \mathbb{R}^N$ $P_{\ker(A)}(x) = x - P_{\operatorname{ran}(A^\top)}(x)$.

$$\ker(A) = \operatorname{ran}(A^\top)^\perp = \big\{ x \in \mathbb{R}^N \;\big|\; Ax = 0 \big\}. \tag{1.30}$$

The associated projection can be derived from $P_{\operatorname{ran}(A^\top)}$ [12, Proposition 3.28(iii)]

$$(\forall x \in \mathbb{R}^N) \qquad P_{\ker(A)}(x) = x - A^\top (A^\top)^\dagger x. \tag{1.31}$$

Finally, let $b \in \operatorname{ran}(A)$, and let

$$C = \big\{ x \in \mathbb{R}^N \;\big|\; Ax = b \big\}. \tag{1.32}$$

The associated projection can be expressed as [12, Example 28.14(ii)]

$$(\forall x \in \mathbb{R}^N) \qquad P_C(x) = x - A^\top (A^\top)^\dagger (x - A^\dagger b). \tag{1.33}$$

In the case when $\operatorname{rank}(A) = K$ (and thus $K \leq N$), it simplifies to

$$(\forall x \in \mathbb{R}^N) \qquad P_C(x) = x - A^\top (AA^\top)^{-1} (Ax - b), \tag{1.34}$$

which, for $K = 1$, further reduces to the projection onto an hyperplane.

Standard simplex and $\boldsymbol{\ell}_1$-ball

Let $\eta \in]0, +\infty[$, and let

$$S = \big\{ x \in [0, +\infty[^N \;\big|\; \sum_{n=1}^N x^{(n)} = \eta \big\}. \tag{1.35}$$

The associated projection can be expressed as [12, Example 28.27]

$$(\forall x \in \mathbb{R}^N) \qquad P_S(x) = \big[\max\{0, x^{(n)} - s\} \big]_{1 \leq n \leq N}, \tag{1.36}$$

where $s \in \mathbb{R}$ is the unique value such that

The sought value of s *can be efficiently computed with the strategies proposed in [71, 217].*

$$s = \sum_{n=1}^{\overline{n}} \frac{\nu^{(n)} - \eta}{\overline{n}} \quad \text{and} \quad \overline{n} = \max \big\{ j \in \{1, \ldots, N\} \;\big|\; \sum_{n=1}^{j} \frac{\nu^{(n)} - \eta}{j} < \nu^{(j)} \big\}, \tag{1.37}$$

with $(\nu^{(n)})_{1 \leq n \leq N}$ being the elements $(x^{(n)})_{1 \leq n \leq N}$ sorted in ascending order. The projection onto the $\boldsymbol{\ell}_1$-ball $C = \big\{ x \in \mathbb{R}^N \;\big|\; \|x\|_1 \leq \eta \big\}$ can be handled similarly, as $P_C(x) = \operatorname{sign}(x) \circ P_S\big(|x|\big)$ for every $x \in \mathbb{R}^N \setminus C$, and $P_C(x) = x$ otherwise, where $\circ$ denotes the element-wise product and $|x| = \big[\, |x^{(n)}| \,\big]_{1 \leq n \leq N}$.

1.3.2 Subgradient projection

The main difficulty in optimizing a convex function over a closed convex set arises when there is no available expression for the corresponding projection. A possible way to circumvent such an issue consists of resorting to the concept of *subgradient projection* [21, 58, 200]. This technique applies to constraints that can be defined as the sublevel set of a function $\varphi \in \Gamma_0(\mathbb{R}^N)$ at level $\eta \in \mathbb{R}$, namely

$$\operatorname{lev}_{\leq \eta} \varphi = \big\{ x \in \mathbb{R}^N \;\big|\; \varphi(x) \leq \eta \big\}. \tag{1.38}$$

The key observation lies in the fact that, for every $u \in \mathbb{R}^N$ and $t \in \partial\varphi(u)$, the above constraint can be outer approximated by the closed half-space

$$C_\varphi^u = \big\{ x \in \mathbb{R}^N \;\big|\; (x-u)^\top t + \varphi(u) \leq \eta \big\}, \tag{1.39}$$

as $x \mapsto (x-u)^\top t + \varphi(u)$ provides a linearization of φ at u, yielding $\operatorname{lev}_{\leq \eta} \varphi \subset C_\varphi^u$ when $t \neq 0$. The projection of u onto C_φ^u is the so-called subgradient projection:

$$P_{C_\varphi^u}(u) = \begin{cases} u, & \text{if } \varphi(u) \leq \eta, \\ u + \dfrac{\eta - \varphi(u)}{\|t\|^2}\, t, & \text{otherwise.} \end{cases} \tag{1.40}$$

The subgradient projection was employed in [58] to optimize a quadratic and strictly-convex function over the sublevel set of a convex function, namely

$$\underset{x \in \mathbb{R}^N}{\text{minimize}} \quad \|Ax - z\|_2^2 \quad \text{s.t.} \quad \varphi(x) \leq \eta, \tag{1.41}$$

where $A \in \mathbb{R}^{K \times N}$ has full column rank, and $z \in \mathbb{R}^K$. The above problem is equivalent to compute the projection of $x^{[0]} = (A^\top A)^{-1} A^\top z$ onto $\operatorname{lev}_{\leq \eta} \varphi$ in the Hilbert space $\mathcal{H}_A$ induced by $\langle \cdot \mid \cdot \rangle_{\mathcal{H}_A} = \langle A^\top A \cdot \mid \cdot \rangle$, since

$$\|Ax - z\|_2^2 = \|x - x^{[0]}\|_{\mathcal{H}_A}^2 - \|x^{[0]}\|_{\mathcal{H}_A}^2 + \|z\|_2^2. \tag{1.42}$$

To tackle this problem, a sequence $(x^{[i]})_{i \in \mathbb{N}}$ is generated by projecting $x^{[0]}$ onto the intersection of two half-spaces built from $C_\varphi^{x^{[i]}}$, yielding the iterations [58]

$$\left\lfloor \begin{aligned} p^{[i]} \quad &= P_{C_\varphi^{x^{[i]}}}(x^{[i]}), \\ x^{[i+1]} &= \underset{x \in \mathcal{H}_A}{\arg\min}\; \|x - x^{[0]}\|_{\mathcal{H}_A}^2 \quad \text{s.t.} \quad \begin{cases} \langle x - x^{[i]} \mid x^{[0]} - x^{[i]} \rangle_{\mathcal{H}_A} \leq 0, \\ \langle x - p^{[i]} \mid x^{[i]} - p^{[i]} \rangle_{\mathcal{H}_A} \leq 0, \end{cases} \end{aligned} \right. \tag{1.43}$$

An essential feature of this method is that the projection onto the intersection of two half-spaces is straightforward to be computed [12, Corollary 28.21]. In addition, an arbitrary number of constraints can be handled in Problem (1.41) by setting $\varphi = \max_{1 \leq r \leq R} \varphi_r$, with $\varphi_r \in \Gamma_0(\mathbb{R}^N)$ for every $r \in \{1, \ldots, R\}$. In this case, a surrogate for the subgradient projection of $\operatorname{lev}_{\leq \eta} \varphi$ is given by the weighted average of the subgradient projections of some $\operatorname{lev}_{\leq \eta} \varphi_r$, with $r \in \mathbb{I}_i \subset \{1, \ldots, R\}$. However, despite this method can handle multiple constraints, it is limited by the fact that the function to be minimized must be quadratic and strictly convex.

1.3.3 Discrepancy principle

An alternative approach to deal with a difficult projection consists of solving the quadratic constrained problem (1.12) through numerical methods [178, Section 6]. This approach may be particularly effective for handling the sublevel set of some function $\varphi \in \Gamma_0(\mathbb{R}^N)$. As a matter of fact, there exists a Lagrangian equivalence between the projection onto $\operatorname{lev}_{\leq \eta} \varphi$ and the proximity operator of φ, in the sense that [12, Proposition 28.30]

$$P_{\operatorname{lev}_{\leq \eta} \varphi}(x) = \begin{cases} x, & \text{if } \varphi(x) \leq \eta, \\ \operatorname{prox}_{\overline{\lambda}\varphi}(x), & \text{otherwise,} \end{cases} \tag{1.44}$$

where $\overline{\lambda}$ is a solution in $]0, +\infty[$ to the equation

$$\varphi\Big(\operatorname{prox}_{\lambda\varphi}(x)\Big) = \eta. \tag{1.45}$$

While there is no general formula for solving the above equality, specific numerical methods can be developed to find $\overline{\lambda}$. However, this approach turns out to be efficient only when the proximity operator admits a simple form, which is the case of the $\boldsymbol{\ell}_1$-norm [71, 217] and the Kullback-Leibler divergence [207]. Note that there also exist numerical methods based on other approaches (e.g., [188]), but they may prove inefficient in large-size problems, due to inner iterations.

1.4 Conclusion

In this chapter, we have presented a quick overview of non-smooth convex optimization. Firstly, we have introduced inverse problems and discussed the advantages of formulating them as constrained convex optimization problems. Then, we have presented some algorithms that can efficiently deal with a large panel of convex problems. Finally, we have highlighted the difficulty of handling nonlinear constraints, we have recalled some examples of projections that can be explicitly computed, and we have made an overview of the existing approaches in the case when the projection admits no closed-form expression. In the next chapter, we will propose a novel technique to deal with a class of nonlinear constraints for which the existing methods are infeasible or inefficient.

Chapter 2
Epigraphical splitting

The main focus of this chapter is convex optimization under nonlinear constraints. Firstly, we propose a technique for decomposing a constraint into some epigraphs. Secondly, we provide closed-form expressions of the projection onto several epigraphs.

2.1 Introduction

Numerous constraints usually involved in the formulation of inverse problems can be modelled through a block-decomposable function. Popular examples are the Kullback-Leibler divergence [207], the hinge loss [78], the $\boldsymbol{\ell}_2$-norm composed with the Anscombe transformation [113], the $\boldsymbol{\ell}_1$-norm [217], the $\boldsymbol{\ell}_{1,\infty}$-norm [188], and total variation [61] or total generalized variation [177] semi-norms. A possible solution to deal with these constraints is to exploit the Lagrangian equivalence between the projection and proximity operators, which boils down to the problem of finding the zero of a nonlinear equation (see Section 1.3.3). However, to the best of our knowledge, this approach turns out to be efficient only with the Kullback-Leibler divergence [207] and the $\boldsymbol{\ell}_1$-norm [71, 217].

The present chapter aims at designing an efficient method for solving the following class of constrained convex optimization problems:

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r x) \quad \text{s.t.} \quad h(Fx) \leq \eta, \tag{2.1}$$

where $T_r \in \mathbb{R}^{K_r \times N}$ and $f_r \in \Gamma_0(\mathbb{R}^{K_r})$ for every $r \in \{1,\dots,R\}$, while $\eta \in \mathbb{R}$, $F \in \mathbb{R}^{M\times N}$, and $h \in \Gamma_0(\mathbb{R}^M)$ is a *separable* function such that

$$(\forall y \in \mathbb{R}^M) \qquad h(y) = \sum_{\ell=1}^{L} h_\ell(\mathrm{y}^{(\ell)}), \tag{2.2}$$

with the generic vector y decomposed into blocks of coordinates as follows

$$y = \Big[\underbrace{\big(\mathrm{y}^{(1)}\big)^\top}_{\text{size } M_1}, \dots, \underbrace{\big(\mathrm{y}^{(L)}\big)^\top}_{\text{size } M_L} \Big]^\top \in \mathbb{R}^{M=M_1+\cdots+M_L}. \tag{2.3}$$

and $h_\ell \in \Gamma_0(\mathbb{R}^{M_\ell})$ such that $\operatorname{ri}(\operatorname{dom} h_\ell) \neq \emptyset$ for every $\ell \in \{1,\dots,L\}$.

Section 2.2 details a technique that splits the constraint in Problem (2.1) into a collection of epigraphs, so as to exchange the projection onto the original constraint set with simpler epigraphical projections. In this regard, Section 2.3 provides novel theoretical results concerning the epigraphical projection of several functions of practical interest, such as the absolute value raised to a power $q \in [1,+\infty[$, the distance to a convex set $C \subset \mathbb{R}^N$, the $\boldsymbol{\ell}_p$-norm with $p \in \{2,+\infty\}$, and the max function. Finally, some conclusions are drawn in Section 2.4.

The epigraph of $\varphi \in \Gamma_0(\mathbb{R}^N)$ is the subset of $\mathbb{R}^N \times \mathbb{R}$ defined as

$$\operatorname{epi}\varphi = \big\{(x,\zeta) \;\big|\; \varphi(x) \leq \zeta\big\}.$$

The distance to a closed convex set $C \subset \mathbb{R}^N$ is defined as

$$d_C = \|\cdot - P_C\|_2.$$

2.2 Proposed method

The *epigraphical splitting* consists of introducing an auxiliary vector $\zeta \in \mathbb{R}^L$ in the minimization process, so that the constraint in Problem (2.1) can be split into a collection of epigraphs and a closed half-space:

The linear inequality over ζ can be also replaced by an equality, even though it makes little difference in our approach.

$$\sum_{\ell=1}^{L} h_\ell(\mathsf{y}^{(\ell)}) \leq \eta \quad \Leftrightarrow \quad \begin{cases} (\forall \ell \in \{1,\dots,L\}) \quad h_\ell(\mathsf{y}^{(\ell)}) \leq \zeta^{(\ell)}, \\ \displaystyle\sum_{\ell=1}^{L} \zeta^{(\ell)} \leq \eta. \end{cases} \tag{2.4}$$

Consequently, Problem (2.1) can be equivalently formulated as follows:

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^L}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\, x) \quad \text{s.t.} \quad \begin{cases} (\forall \ell \in \{1,\dots,L\}) \quad h_\ell\left(F_\ell\, x\right) \leq \zeta^{(\ell)}, \\ \displaystyle\sum_{\ell=1}^{L} \zeta^{(\ell)} \leq \eta, \end{cases} \tag{2.5}$$

where, for every $\ell \in \{1,\dots,L\}$, $F_\ell \in \mathbb{R}^{M_\ell \times N}$ is such that $F_\ell\, x = [Fx]^{(\ell)} = \mathsf{y}^{(\ell)}$. Note that the above minimization problem is defined with respect to x and ζ, so we have increased the dimensionality of our problem, and we have replaced the sublevel set of h with simpler constraints given by the epigraphs of $h_1,\dots,h_L$.

2.2.1 Connections with proximal algorithms

Within the proposed constrained optimization framework, Problem (2.5) can be rewritten in a more compact form as follows:

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^M}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r x) \quad \text{s.t.} \quad \begin{cases} (Fx,\zeta) \in E, \\ \zeta \in V, \end{cases} \tag{2.6}$$

where

$$E = \left\{(y,\zeta) \in \mathbb{R}^M \times \mathbb{R}^L \;\middle|\; (\forall \ell \in \{1,\dots,L\}) \quad (\mathsf{y}^{(\ell)},\zeta^{(\ell)}) \in \operatorname{epi} h_\ell\right\}, \tag{2.7}$$

$$V = \left\{\zeta \in \mathbb{R}^L \;\middle|\; \mathsf{1}_L^\top \zeta \leq \eta\right\}, \tag{2.8}$$

with $\mathsf{1}_L = (1,\dots,1)^\top \in \mathbb{R}^L$. The above problem can be efficiently solved by the primal-dual proximal methods presented in Section 1.2.4, provided that the operators $(\operatorname{prox}_{f_r})_{1\leq r\leq R}$, P_E, and P_V can be quickly computed. In the present case, we assume that $(\operatorname{prox}_{f_r})_{1\leq r\leq R}$ have closed-form expressions. In addition, the projection onto V is given in Section 1.3.1, whereas

$$(\forall (y,\zeta) \in \mathbb{R}^M \times \mathbb{R}^L) \qquad P_E(y,\zeta) = (p,\theta), \tag{2.9}$$

where $p \in \mathbb{R}^M$ is blockwise decomposed as in (2.3), $\theta \in \mathbb{R}^L$, and

$$(\forall \ell \in \{1,\dots,L\}) \qquad (\mathsf{p}^{(\ell)},\theta^{(\ell)}) = P_{\operatorname{epi} h_\ell}(\mathsf{y}^{(\ell)},\zeta^{(\ell)}). \tag{2.10}$$

Consequently, in order to solve Problem (2.6), we need to compute the projection onto $\operatorname{epi} h_\ell$ for each $\ell \in \{1,\dots,L\}$, which yields two potential benefits with respect to Problem (2.1). Firstly, the epigraphical projection involves the lower-dimensional problem of determining the projection onto the convex subset $\operatorname{epi} h_\ell$ of $\mathbb{R}^{M_\ell} \times \mathbb{R}$. Secondly, these projections can be computed in parallel, since they are defined over disjoint blocks [102]. Examples of algorithms that solve Problem (2.6), and thus Problem (2.1), will be provided in the next chapters.

2.2.2 Examples of epigraphical constraints

We now illustrate some examples of functions that can be handled with the epigraphical splitting presented above. The mathematical expression of the associated projections will be derived in Section 2.3.

(i). **$\boldsymbol{\ell}_q$-norm**. Let $q \geq 1$, and let $\tau_\ell > 0$ for every $\ell \in \{1, \ldots, L\}$. Then, the function

$$(\forall y \in \mathbb{R}^M) \qquad h(y) = \sum_{\ell=1}^{M} \tau_\ell \, |y^{(\ell)}|^q \tag{2.11}$$

can be modelled as in (2.2) with $L = M$ and

$$(\forall \ell \in \{1, \ldots, M\})(\forall y^{(\ell)} \in \mathbb{R}) \qquad h_\ell(y^{(\ell)}) = \tau_\ell \, |y^{(\ell)}|^q. \tag{2.12}$$

The corresponding epigraphical projection will be given in Proposition 2.3.2 for $q = 1$ and Proposition 2.3.3 for $q > 1$. Note that the $\boldsymbol{\ell}_1$ and $\boldsymbol{\ell}_2$ norms are widely used for the regularization of inverse problems [34, 209].

(ii). **Distance function**. For every $\ell \in \{1, \ldots, L\}$, let $\tau_\ell > 0$, $q_\ell \geq 1$, and $C_\ell \subset \mathbb{R}^{M_\ell}$ be a nonempty closed convex set. Then, the function

$$(\forall y \in \mathbb{R}^M) \qquad h(y) = \sum_{\ell=1}^{L} \tau_\ell \, d_{C_\ell}^{q_\ell}(\mathsf{y}^{(\ell)}) \tag{2.13}$$

can be modelled as in (2.2) with

$$(\forall \ell \in \{1, \ldots, L\})(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad h_\ell(\mathsf{y}^{(\ell)}) = \tau_\ell \, d_{C_\ell}^{q_\ell}(\mathsf{y}^{(\ell)}). \tag{2.14}$$

The related epigraphical projection will be given in Proposition 2.3.4. Such a function is relevant for relaxing constraints on support sets [64].

(iii). **$\boldsymbol{\ell}_{1,2}$-norm**. Let $\tau_\ell > 0$ for every $\ell \in \{1, \ldots, L\}$. Then, the function

$$(\forall y \in \mathbb{R}^M) \qquad h(y) = \sum_{\ell=1}^{L} \tau_\ell \, \|\mathsf{y}^{(\ell)}\|_2 \tag{2.15}$$

can be modelled as in (2.2) with

$$(\forall \ell \in \{1, \ldots, L\})(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad h_\ell(\mathsf{y}^{(\ell)}) = \tau_\ell \, \|\mathsf{y}^{(\ell)}\|_2. \tag{2.16}$$

The associated epigraphical projection will be given in Corollary 2.3.5. Note that the $\boldsymbol{\ell}_{1,2}$-norm is useful to define multivariate sparsity constraints [80, 226] or total variation bounds [195], which typically involve a sum of functions like (2.16) composed with frames or gradient operators.

(iv). **$\boldsymbol{\ell}_{1,\infty}$-norm**. For every $\ell \in \{1, \ldots, L\}$, let $\Omega_\ell = \mathrm{Diag}(\tau_{\ell,1}, \ldots, \tau_{\ell,M_\ell})$, with $\tau_{\ell,m} > 0$ for each $m \in \{1, \ldots, M_\ell\}$. Then, the function

$$(\forall y \in \mathbb{R}^M) \qquad h(y) = \sum_{\ell=1}^{L} \|\Omega_\ell \, \mathsf{y}^{(\ell)}\|_\infty \tag{2.17}$$

can be modelled as in (2.2) with

$$(\forall \ell \in \{1,\dots,L\})(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad h_\ell(\mathsf{y}^{(\ell)}) = \max_{1\leq m\leq M_\ell} \tau_{\ell,m}\,|y^{(\ell,m)}|, \tag{2.18}$$

where $\mathsf{y}^{(\ell)} = (y^{(\ell,m)})_{1\leq m\leq M_\ell}$. The corresponding epigraphical projection will be given in Proposition 2.3.7. The infinity norm $\|\cdot\|_\infty$ has recently attracted much interest for regularization purposes [42, 188, 204].

(v). **Max function**. Let $r_{\ell,m} \in \mathbb{R}$ for every $\ell \in \{1,\dots,L\}$ and $m \in \{1,\dots,M_\ell\}$. Then, the function

$$(\forall y \in \mathbb{R}^M) \qquad h(y) = \sum_{\ell=1}^{L} \max_{1\leq m\leq M_\ell} y^{(\ell,m)} + r_{\ell,m} \tag{2.19}$$

can be modelled as in (2.2) with

$$(\forall \ell \in \{1,\dots,L\})(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad h_\ell(\mathsf{y}^{(\ell)}) = \max_{1\leq m\leq M_\ell} y^{(\ell,m)} + r_{\ell,m}, \tag{2.20}$$

where $\mathsf{y}^{(\ell)} = (y^{(\ell,m)})_{1\leq m\leq M_\ell}$. The corresponding epigraphical projection will be given in Proposition 2.3.8. Note that the max function can be used to define the hinge loss involved in multiclass support vector machines [78].

2.2.3 Differences with existing splitting techniques

Splitting methods play a central role in non-smooth convex optimization. In this context, a number of techniques have been developed to reduce the complexity related to a non-smooth function composed with a linear operator. Aside from the Fenchel-Rockafellar duality (see Section 1.2.3), a popular approach to circumvent this issue is inspired by the Alternating Direction Method of Multipliers [91], which consists of dealing with optimization problems of the form

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad f(x) + h(Fx) \tag{2.21}$$

by resorting to the following (equivalent) reformulation

$$\underset{(x,y)\in\mathbb{R}^N\times\mathbb{R}^M}{\text{minimize}} \quad f(x) + h(y) \quad \text{s.t.} \quad Fx = y. \tag{2.22}$$

This type of splitting has been used in image restoration [1] and, more recently, for distributed optimization problems [23]. A similar approach was also used to derive extensions of Douglas-Rachford method [30, 180], since the constraint $Fx = y$ can be interpreted either as the kernel of $[F \;\; -\mathrm{Id}]$ or the range of $[\mathrm{Id} \;\; F^\top]^\top$, whose projections can be efficiently computed when F is a circulant matrix (see Section 1.3.1 for the definitions of kernel and range).

The solution that we propose in this work also introduces auxiliary variables. However, our objective is not to deal with linear operators, but rather with a projection that does not have a closed-form expression. Consequently, the proposed solution departs from the usual splitting methods, in the sense that our approach leads to a collection of epigraphs (i.e., nonlinear constraints), while the usual splitting techniques involve affine sets (for which the projection is known).

2.3 Epigraphical projections

The key point in the proposed splitting is the introduction of some epigraphs in the minimization process, in order to facilitate the computational steps. Therefore, it is of paramount importance that the projection onto the epigraph can be efficiently computed. The problem of determining such an *epigraphical projection* is formalized in the following proposition.

Proposition 2.3.1. *Let $\mathcal{H}$ be a real Hilbert space and let $\mathcal{H}\times\mathbb{R}$ be equipped with the standard product space norm. Let φ be a function in $\Gamma_0(\mathcal{H})$ such that* $\operatorname{dom}\varphi$ *is open. For every $(y,\zeta)\in\mathcal{H}\times\mathbb{R}$, the projection onto* $\operatorname{epi}\varphi$ *is given by*

$$P_{\operatorname{epi}\varphi}(y,\zeta)=\Big(p,\ \max\{\varphi(p),\zeta\}\Big), \tag{2.23}$$

where

$$p=\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y). \tag{2.24}$$

Proof. For every $(y,\zeta)\in\mathcal{H}\times\mathbb{R}$, $P_{\operatorname{epi}\varphi}(y,\zeta)=(p,\theta)$ is the unique solution to

$$\underset{(u,\xi)\in\mathcal{H}\times\mathbb{R}}{\text{minimize}}\quad \|u-y\|^2+(\xi-\zeta)^2\quad \text{s.t.}\quad \varphi(u)\le\zeta. \tag{2.25}$$

If $\varphi(y)\le\zeta$, the solution is $(p,\theta)=(y,\zeta)$, which is equivalent to (2.23) because $\varphi(y)\le\zeta$ implies $\max\{\varphi(y)-\zeta,0\}=0$, and thus $\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y)=y$ and $\max\{\varphi(y),\zeta\}=\zeta$. On the other side, if $\varphi(y)>\zeta$, the Karush-Kuhn-Tucker theorem [12, Proposition 26.18] assures the existence of $\alpha\in[0,+\infty[$ such that

$$(p,\theta)=\underset{(u,\xi)\in\mathcal{H}\times\mathbb{R}}{\operatorname{argmin}}\ \frac{1}{2}\|u-y\|^2+\frac{1}{2}(\xi-\zeta)^2+\alpha(\varphi(u)-\xi), \tag{2.26}$$

The required qualification condition is obviously satisfied when $u\in\operatorname{dom}\varphi$ and $\xi\ge\varphi(u)$

where α is such that $\alpha(\varphi(p)-\theta)=0$. Since the value $\alpha=0$ is not allowable (otherwise $p=y$ and $\theta=\zeta$), we deduce that $\varphi(p)=\theta$ and $\alpha>0$. Moreover, differentiating the Lagrange functional in (2.26) yields $y-p\in\alpha\,\partial\varphi(p)$ and $\theta=\zeta+\alpha$. Since $\alpha>0$, it results $\theta>\zeta$, and thus $\theta=\max\{\theta,\zeta\}=\max\{\varphi(p),\zeta\}$, as $\theta=\varphi(p)$. In addition, the solution to (2.25) is $(p,\theta)=\big(p,\varphi(p)\big)$ with

$$p=\underset{\substack{u\in\mathcal{H}\\ \varphi(u)\ge\zeta}}{\operatorname{argmin}}\ \|u-y\|^2+(\varphi(u)-\zeta)^2. \tag{2.27}$$

Furthermore, as $\varphi(y)>\zeta$, we obtain

$$\inf_{\substack{u\in\mathcal{H}\\ \varphi(u)\le\zeta}}\|u-y\|^2=\|P_{\operatorname{lev}_{\le\zeta}\varphi}(y)-y\|^2=\inf_{\substack{u\in\mathcal{H}\\ \varphi(u)=\zeta}}\|u-y\|^2, \tag{2.28}$$

where we have used the fact that $P_{\operatorname{lev}_{\le\zeta}\varphi}(y)$ belongs to $\big\{u\in\mathcal{H}\ \big|\ \varphi(u)=\zeta\big\}$, since $\varphi\in\Gamma_0(\mathcal{H})$ and $\operatorname{dom}\varphi$ is open [12, Corollary 8.38]. We have then

$$\inf_{\substack{u\in\mathcal{H}\\ \varphi(u)\le\zeta}}\|u-y\|^2=\inf_{\substack{u\in\mathcal{H}\\ \varphi(u)=\zeta}}\|u-y\|^2\ge\inf_{\substack{u\in\mathcal{H}\\ \varphi(u)\ge\zeta}}\|u-y\|^2+(\varphi(u)-\zeta)^2. \tag{2.29}$$

Altogether, (2.27) and (2.29) lead to

$$p=\underset{u\in\mathcal{H}}{\operatorname{argmin}}\ \frac{1}{2}\|u-y\|^2+\frac{1}{2}\big(\varphi(u)-\zeta\big)^2 \tag{2.30}$$

which is equivalent to (2.24) as $\frac{1}{2}(\max\{\varphi-\zeta,0\})^2\in\Gamma_0(\mathcal{H})$. □

2.3.1 Absolute value

The previous proposition shows that the proximity operator in (2.24) plays a prominent role in computing the projection onto an epigraph. Interestingly, such a proximity operator admits a simple form for several functions of practical interest, such as the absolute value, as shown by the next propositions.

Proposition 2.3.2. *Let $\tau \in]0, +\infty[$. Assume that*

$$(\forall y \in \mathbb{R}) \qquad \varphi(y) = \tau|y|. \tag{2.31}$$

For every $(y, \zeta) \in \mathbb{R} \times \mathbb{R}$, the projection onto $\operatorname{epi}\varphi$ *is given by* (2.23) *with*

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} y, & \text{if } \tau|y| \leq \zeta, \\ \dfrac{\operatorname{sign}(y)}{1+\tau^2} \max\{|y| + \tau\zeta, 0\}, & \text{otherwise.} \end{cases} \tag{2.32}$$

Proof. The above result follows by the fact that, for every $(y, \zeta) \in \mathbb{R}^2$,

$$\frac{1}{2}(\max\{\tau|y| - \zeta, 0\})^2 = \begin{cases} 0, & \text{if } \tau|y| \leq \zeta, \\ (\tau^2/2)\, y^2 - \tau\zeta|y| + \zeta^2/2, & \text{otherwise,} \end{cases} \tag{2.33}$$

for which the proximity operator is known [40, Example 4.6]. □

Proposition 2.3.3. *Let $q \in]1, +\infty[$ and $\tau \in]0, +\infty[$. Assume that*

$$(\forall y \in \mathbb{R}) \qquad \varphi(y) = \tau|y|^q. \tag{2.34}$$

For every $(y, \zeta) \in \mathbb{R} \times \mathbb{R}$, the projection onto $\operatorname{epi}\varphi$ *is given by* (2.23) *with*

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} y, & \text{if } \tau|y|^q \leq \zeta, \\ \operatorname{sign}(y)\, \chi, & \text{otherwise,} \end{cases} \tag{2.35}$$

When q is a rational number, (2.36) *is equivalent to a polynomial equation for which there exists a closed-form solution or a standard numerical method*

where χ is the unique solution on $[(\max\{\zeta, 0\}/\tau)^{1/q}, +\infty[$ of the equation

$$q\tau^2\chi^{2q-1} - q\tau\zeta\chi^{q-1} + \chi = |y|. \tag{2.36}$$

Proof. Since $(\max\{\tau|\cdot|^q - \zeta, 0\})^2$ is an even function, $\operatorname{prox}_{\frac{1}{2}(\max\{\tau|\cdot|^q-\zeta,0\})^2}$ is an odd function [40, Remark 4.1(ii)]. In the following, we thus focus on the case when $y > 0$. If $\zeta \leq 0$, then $(\max\{\tau(\cdot)^q - \zeta, 0\})^2 = (\tau(\cdot)^q - \zeta)^2$ is differentiable, thus we can deduce that $\chi = \operatorname{prox}_{\frac{1}{2}(\tau(\cdot)^q-\zeta)^2}(y)$ is equivalent to

$$y - \chi = \nabla\left(\frac{1}{2}(\tau\chi^q - \zeta)^2\right) \quad \Leftrightarrow \quad \chi - y + q\tau\chi^{q-1}(\tau\chi^q - \zeta) = 0, \tag{2.37}$$

where $\chi \geq 0$ by virtue of [63, Corollary 2.5]. This allows us to state that χ is the unique solution of (2.36) on $[0, +\infty[$. Let us now focus on the case when $\zeta > 0$. If $0 < y \leq (\zeta/\tau)^{1/q}$, then $\max\{\varphi - \zeta, 0\} = 0$, thus $\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = y$. On the other hand, if $y > (\zeta/\tau)^{1/q}$, as the proximity operator of a function from $\mathbb{R}$ to $\mathbb{R}$ is continuous and increasing [63, Proposition 2.4], we obtain

$$\chi = \operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) \geq \operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}\left((\zeta/\tau)^{1/q}\right) = (\zeta/\tau)^{1/q}. \tag{2.38}$$

Therefore, χ is the unique solution to (2.36) on $[(\zeta/\tau)^{1/q}, +\infty[$. □

2.3.2 Distance function

The previous propositions allow us to establish a result concerning the distance function to a convex set $C \subset \mathcal{H}$. Note that, in the case when $C = \{z\}$ for some $z \in \mathcal{H}$, this function reduces to $d_C = \|\cdot - z\|_2$, as $P_C = z$. Hence, the next proposition leads to a corollary about the Euclidean norm, for which the associated epigraph is called the *Lorentz cone* and its epigraphical projection is known in the literature (e.g., see [178, Section 6.3.2]).

Proposition 2.3.4. *Let C be a nonempty convex subset of $\mathcal{H}$. Let $q \in [1, +\infty[$, $\tau \in]0, +\infty[$, and $\zeta \in \mathbb{R}$. Assume that*

$$(\forall y \in \mathcal{H}) \qquad \varphi(y) = \tau d_C^q(y). \tag{2.39}$$

For every $(y, \zeta) \in \mathcal{H} \times \mathbb{R}$, the projection onto $\operatorname{epi}\varphi$ is given by (2.23) with

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} y, & \text{if } y \in C, \\ \alpha y + (1-\alpha)P_C(y), & \text{otherwise,} \end{cases} \tag{2.40}$$

where

$$\alpha = \frac{\operatorname{prox}_{\frac{1}{2}(\max\{\tau|\cdot|^q-\zeta,0\})^2}\big(d_C(y)\big)}{d_C(y)}, \tag{2.41}$$

and the above expression is provided by Propositions 2.3.2–2.3.3.

Proof. Note that $\frac{1}{2}(\max\{\tau d_C^q - \zeta, 0\})^2 = \psi \circ d_C$, with $\psi = \frac{1}{2}(\max\{\tau|\cdot|^q - \zeta, 0\})^2$. According to [64, Proposition 2.7], for every $y \in \mathcal{H}$,

$$\operatorname{prox}_{\psi \circ d_C}(y) = \begin{cases} y, & \text{if } y \in C, \\ P_C(y), & \text{if } d_C(y) \leq \max \partial\psi(0), \\ \alpha y + (1-\alpha)P_C(y), & \text{if } d_C(y) > \max \partial\psi(0), \end{cases} \tag{2.42}$$

where $\alpha = \operatorname{prox}_{\psi}\big(d_C(y)\big)/d_C(y)$. In addition, we have $\partial\psi(0) = [\tau\zeta, -\tau\zeta]$ if $\zeta < 0$ and $q = 1$, or $\partial\psi(0) = \{0\}$ otherwise. As a result, (2.42) translates to

$$\operatorname{prox}_{\psi \circ d_C}(y) = \begin{cases} y, & \text{if } y \in C, \\ P_C(y), & \text{if } d_C(y) \leq -\tau\zeta \text{ and } q = 1, \\ \alpha y + (1-\alpha)P_C(y), & \text{otherwise.} \end{cases} \tag{2.43}$$

By Proposition 2.3.2, $\operatorname{prox}_{\psi}\big(d_C(y)\big) = 0$ when $d_C(y) \leq -\tau\zeta$ and $q = 1$, which implies $\alpha = 0$ and $P_C(y) = \alpha y + (1-\alpha)P_C(y)$. Hence, (2.43) yields (2.40). □

Corollary 2.3.5. *Let $\tau \in]0, +\infty[$, $\zeta \in \mathbb{R}$ and $z \in \mathcal{H}$. Assume that*

$$(\forall y \in \mathcal{H}) \qquad \varphi(y) = \tau \|y - z\|_2. \tag{2.44}$$

For every $(y, \zeta) \in \mathcal{H} \times \mathbb{R}$, the projection onto $\operatorname{epi}\varphi$ is given by (2.23) with

$$\operatorname{prox}_{\frac{1}{2}(\max\{\varphi-\zeta,0\})^2}(y) = \begin{cases} z, & \text{if } y = z, \\ y, & \text{if } \tau\|y - z\|_2 \leq \zeta, \\ z + \alpha(y - z), & \text{otherwise,} \end{cases} \tag{2.45}$$

where

$$\alpha = \frac{1}{1+\tau^2} \max\left\{1 + \frac{\tau\zeta}{\|y - z\|_2}, 0\right\}. \tag{2.46}$$

2.3.3 Max function

We conclude the section with some results about the epigraphical projection of the max function. The first one considers the maximum of weighted absolute values (which reduces to the standard infinity norm $\|\cdot\|_\infty$ when the weights are all equal [85]), whereas the second one concerns the maximum of weighted values. The proofs of these propositions are very similar, and rely on the expression of the proximity operator involved in the next lemma.

Lemma 2.3.6. *Let $(\tau_m)_{1\leq m\leq M} \in \mathbb{R}^M$ and $\nu = (\nu^{(m)})_{1\leq m\leq M} \in \mathbb{R}^M$. Assume that*

$$(\forall y \in \mathbb{R}) \qquad \varphi(y) = \frac{1}{2}\sum_{m=1}^{M} \big(\max\{\tau_m\,(\nu^{(m)} - y), 0\}\big)^2, \tag{2.47}$$

with the values $(\nu^{(m)})_{1\leq m\leq M_\ell}$ sorted in ascending order. Then, $\varphi \in \Gamma_0(\mathbb{R})$ and

$$(\forall y \in \mathbb{R}) \qquad \operatorname{prox}_{\varphi}(y) = \frac{y + \sum_{m=1}^{\overline{m}-1} \nu^{(m)}(\tau_m^-)^2 + \sum_{m=\overline{m}}^{M} \nu^{(m)}(\tau_m^+)^2}{1 + \sum_{m=1}^{\overline{m}-1}(\tau_m^-)^2 + \sum_{m=\overline{m}}^{M}(\tau_m^+)^2}, \tag{2.48}$$

where $\tau_m^- = \min\{\tau_m, 0\}$ and $\tau_m^+ = \max\{\tau_m, 0\}$ for every $m \in \{1,\dots,M\}$, whereas $\overline{m}$ is the unique integer in $\{1,\dots,M+1\}$ such that

$$\nu^{(\overline{m}-1)} < \frac{y + \sum_{m=1}^{\overline{m}-1} \nu^{(m)}(\tau_m^-)^2 + \sum_{m=\overline{m}}^{M} \nu^{(m)}(\tau_m^+)^2}{1 + \sum_{m=1}^{\overline{m}-1}(\tau_m^-)^2 + \sum_{m=\overline{m}}^{M}(\tau_m^+)^2} \leq \nu^{(\overline{m})}, \tag{2.49}$$

with the conventions $\nu^{(0)} = -\infty$, $\nu^{(M+1)} = +\infty$, and $\sum_{m=1}^{0} \cdot = \sum_{m=M+1}^{M} \cdot = 0$.

Proof. The function φ belongs to $\Gamma_0(\mathbb{R})$, as $\max\{\tau_m\,(\nu^{(m)} - \cdot), 0\}$ is finite convex, and $(\cdot)^2$ is finite convex and increasing on $[0, +\infty[$. In addition, φ is differentiable and such that, for every $v \in \mathbb{R}$ and $k \in \{1,\dots,M+1\}$,

$$\begin{aligned} \nu^{(k-1)} < v \leq \nu^{(k)} \quad \Rightarrow \quad \varphi(v) &= \frac{1}{2}\sum_{m=1}^{k-1}(\tau_m^-)^2\big(v - \nu^{(m)}\big)^2 \\ &\quad + \frac{1}{2}\sum_{m=k}^{M}(\tau_m^+)^2\big(v - \nu^{(m)}\big)^2. \end{aligned} \tag{2.50}$$

Since $p = \operatorname{prox}_{\varphi}(y)$ is uniquely defined by $y - p = \varphi'(p)$, there exists a unique $\overline{m} \in \{1,\dots,M+1\}$ such that $\nu^{(\overline{m}-1)} < p \leq \nu^{(\overline{m})}$ and

$$y - p = \sum_{m=1}^{\overline{m}-1}(\tau_m^-)^2(p - \nu^{(m)}) + \sum_{m=\overline{m}}^{M}(\tau_m^+)^2(p - \nu^{(m)}). \tag{2.51}$$

This yields (2.48), and thus $\nu^{(\overline{m}-1)} < p \leq \nu^{(\overline{m})}$ is equivalent to (2.49). □

Proposition 2.3.7. *Let* $(\tau_m)_{1\leq m\leq M} \in]0,+\infty[^M$. *Assume that*

$$(\forall y \in \mathbb{R}^M) \qquad \varphi(y) = \max_{1\leq m\leq M} \tau_m \, |y^{(m)}|, \tag{2.52}$$

with the values $\big(\nu^{(m)} = \tau_m \, |y^{(m)}|\big)_{1\leq m\leq M}$ *being sorted in ascending order. For every* $(y,\zeta) \in \mathbb{R}^M \times \mathbb{R}$, *the projection* $P_{\mathrm{epi}\,\varphi}(y,\zeta) = (p,\theta)$ *is given by*

$$p = \Big[\operatorname{sign}(y^{(m)}) \, \min\big\{|y^{(m)}|, \theta/\tau_m\big\}\Big]_{1\leq m\leq M} \tag{2.53}$$

and

$$\theta = \max\left\{0, \; \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^{-2}\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^{-2} \, \nu^{(m)}\right)\right\}, \tag{2.54}$$

where $\overline{m}$ *is the unique integer in* $\{1,\dots,M+1\}$ *such that*

$$\nu^{(\overline{m}-1)} < \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^{-2}\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^{-2} \, \nu^{(m)}\right) \leq \nu^{(\overline{m})}, \tag{2.55}$$

with the conventions $\nu^{(0)} = -\infty$, $\nu^{(M+1)} = +\infty$, *and* $\sum_{m=M+1}^{M} \cdot = 0$.

Proof. For every $(y,\zeta) \in \mathbb{R}^M \times \mathbb{R}$, $P_{\mathrm{epi}\,\varphi}(y,\zeta) = (p,\theta)$ is the unique solution to

$$\begin{array}{ll} \underset{\theta\in[0,+\infty[}{\text{minimize}} & \underset{\substack{|p^{(1)}| \leq \theta/\tau_1 \\ \vdots \\ |p^{(M)}| \leq \theta/\tau_M}}{\text{minimize}} \quad (\theta-\zeta)^2 + \|p-y\|^2. \end{array} \tag{2.56}$$

For every $\theta \in [0,+\infty[$, the inner minimization is achieved when, for every $m \in \{1,\dots,M\}$, $p^{(m)}$ is the projection of $y^{(m)}$ onto the range $[-\theta/\tau_m, \theta/\tau_m]$, which is given by (2.54). Then, (2.56) reduces to

$$\underset{\theta\in[0,+\infty[}{\text{minimize}} \quad (\theta-\zeta)^2 + \sum_{m=1}^{M} (\max\{|y^{(m)}| - \theta/\tau_m, 0\})^2, \tag{2.57}$$

which yields $\theta = \operatorname{prox}_{\psi+\iota_{[0,+\infty[}}(\zeta)$ with $\psi = \sum_{m=1}^{M} (\max\{\tau_m^{-1}(\nu^{(m)} - \cdot), 0\})^2$. Hence, we can deduce from [62, Proposition 12] that $\theta = P_{[0,+\infty[}(\chi)$ with $\chi = \operatorname{prox}_{\phi}(\zeta)$, which leads to (2.55) according to Lemma 2.3.6. □

Proposition 2.3.8. *Let* $(r_m)_{1\leq m\leq M} \in \mathbb{R}^M$, $(\tau_m)_{1\leq m\leq M} \in]0,+\infty[^M$, *and* $(\epsilon_m)_{1\leq m\leq M} \in \{1,-1\}^M$. *Assume that*

$$(\forall y \in \mathbb{R}^M) \qquad \varphi(y) = \max_{1\leq m\leq M} \frac{\epsilon_m\, y^{(m)} + r_m}{\tau_m}, \tag{2.58}$$

with the values $\big(\nu^{(m)} = \frac{\epsilon_m y^{(m)}+r_m}{\tau_m}\big)_{1\leq m\leq M}$ *being sorted in ascending order. For every* $(y,\zeta) \in \mathbb{R}^M \times \mathbb{R}$, *the projection* $P_{\mathrm{epi}\,\varphi}(y,\zeta) = (p,\theta)$ *is given by*

$$p = \Big[\epsilon_m \min\big\{\epsilon_m\, y^{(m)}, \tau_m\,\theta - r_m\big\}\Big]_{1\leq m\leq M}, \tag{2.59}$$

and

$$\theta = \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^2\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^2\, \nu^{(m)}\right), \tag{2.60}$$

where $\overline{m}$ *is the unique integer in* $\{1,\ldots,M+1\}$ *such that*

$$\nu^{(\overline{m}-1)} < \left(1 + \sum_{m=\overline{m}}^{M} \tau_m^2\right)^{-1} \left(\zeta + \sum_{m=\overline{m}}^{M} \tau_m^2\, \nu^{(m)}\right) \leq \nu^{(\overline{m})}, \tag{2.61}$$

with the conventions $\nu^{(0)} = -\infty$, $\nu^{(M+1)} = +\infty$, *and* $\sum_{m=M+1}^{M} \cdot = 0$.

Proof. For every $(y,\zeta) \in \mathbb{R}^M \times \mathbb{R}$, $P_{\mathrm{epi}\,\varphi}(y,\zeta) = (p,\theta)$ is the unique solution to

$$\begin{array}{ll} \underset{\theta^{(\ell)}\in\mathbb{R}}{\text{minimize}} & \underset{\substack{\epsilon_1 p^{(1)} \leq \tau_1\theta - r_1 \\ \vdots \\ \epsilon_M p^{(M)} \leq \tau_M\theta - r_M}}{\text{minimize}} \quad (\theta-\zeta)^2 + \|p-y\|^2. \end{array} \tag{2.62}$$

For every $\theta \in \mathbb{R}$ and $m \in \{1,\ldots,M\}$, the inner minimization is achieved when $p^{(m)}$ is the projection of $y^{(m)}$ onto $]-\infty, \tau_m\theta - r_m]$ or $[-(\tau_m\theta - r_m), +\infty[$ according to whether $\epsilon_m = 1$ or $\epsilon_m = -1$, respectively. These projections are given by (2.59). Then, (2.62) reduces to

$$\underset{\theta^{(\ell)}\in\mathbb{R}}{\text{minimize}} \quad (\theta-\zeta)^2 + \sum_{m=1}^{M} \big(\max\{\epsilon_m\, y^{(m)} + r_m - \tau_m\theta, 0\}\big)^2, \tag{2.63}$$

which is equivalent to calculate $\theta = \mathrm{prox}_{\psi}(\zeta)$ with

$$(\forall u \in \mathbb{R}) \qquad \psi(u) = \frac{1}{2}\sum_{m=1}^{M} (\max\{\tau_m(\nu^{(m)} - u), 0\})^2. \tag{2.64}$$

By using Lemma 2.3.6, we can deduce the expressions of (2.60) and (2.61). □

2.4 Conclusion

In this chapter, we have proposed a new approach to deal with a class of constrained convex optimization problems. Firstly, we have introduced a splitting technique to decompose the sublevel set of a block-decomposable function into a collection of epigraphs. Secondly, we have proved a relation between the projection onto an epigraph and the proximity operator of a composite function, from which we have derived the closed-form expressions associated to the epigraphical projection of the absolute value, the distance to a convex set, the Euclidean norm, the infinity norm, and the max function.

In the next chapter, we will turn our attention to mixed-norm constraints. In particular, we will show through an application in image recovery that regularity constraints based on the nonlocal total variation penalty can be efficiently handled by the proposed epigraphical splitting.

A fact is a simple statement that everyone believes. It is innocent, unless found guilty. A hypothesis is a novel suggestion that no one wants to believe. It is guilty, until found effective.

EDWARD TELLER

Chapter 3
CONSTRAINTS BASED ON MIXED NORMS

The main focus of this chapter is the epigraphical splitting applied to mixed-norm constraints. We show through an example of image recovery that regularity constraints based on nonlocal total variation can be efficiently handled by the proposed approach.

3.1 IMAGE RECOVERY PROBLEMS

Natural images are often degraded by blur and noise arising from sensor imprecisions or physical limitations. The goal of *image recovery* is to restore the visual content of a corrupted image by inverting the corresponding degradation process. In this context, a popular task consists of recovering an image $\overline{x} \in \mathbb{R}^N$ as close as possible to an observation $z \in \mathbb{R}^K$ generated through the linear model

$$z = A\overline{x} + b, \tag{3.1}$$

where $A \in \mathbb{R}^{K \times N}$ (with $K \leq N$) describes the physical laws linking $\overline{x}$ to the observed data, and $b \in \mathbb{R}^K$ is a realization of zero-mean white Gaussian noise.

A usual approach to recover $\overline{x}$ from z is to follow a variational approach that aims at solving the convex optimization problem expressed as

$$\underset{x \in [0,255]^N}{\text{minimize}} \quad \|Ax - z\|_2^2 \quad \text{s.t.} \quad h(Fx) \leq \eta, \tag{3.2}$$

where $\eta \geq 0$, $F \in \mathbb{R}^{M \times N}$ (with $M \geq N$) is the linear operator associated with an analysis transformation, and $h \in \Gamma_0(\mathbb{R}^M)$. Hereabove, the quadratic term aims at insuring that the solution is close to the observation, while the constraint imposes some regularity on the solution. Regularization is essential for solving such an ill-posed problem, as it allows one to convey some prior knowledge about the image to be recovered, independently from any specific information that can be inferred from the observed data. In this regard, the more carefully the regularity is modeled, the better the quality of the estimated image.

Natural images usually exhibit a smooth spatial behaviour, except around some locations (such as object edges), where discontinuities arise. Therefore, the quality of the results obtained through the resolution of Problem (3.2) strongly depends on the ability of the operator F and the function h to model such a specific type of regularity. Among the sophisticated regularization terms developed in the field of image restoration (e.g., see [105, 124]), the most popular ones are expressed as the mixed norm of some transformed coefficients, yielding

$$(\forall y \in \mathbb{R}^M) \qquad h(y) = \sum_{\ell=1}^{L} \|\mathsf{y}^{(\ell)}\|_p^q, \tag{3.3}$$

where $q \geq 1$, $p \geq 1$, and y is block-decomposed as in (2.3). Various forms of regularization arise with specific choices of F, p and q, as discussed next.

3.1.1 Quadratic regularization

The choice $p = q = 2$ in (3.3) yields a quadratic penalty, which is one of the most widely referenced approaches for the regularization of ill-posed problems [209]. When F is an orthonormal transformation, the term $h(Fx) = \|Fx\|_2^2$ simply measures the energy of x, and thus the regularity constraint in (3.2) directly prevents the solution coefficients from becoming too large. This regularization is known to induce a *low-pass filtered* solution, meaning that no attempt is made to recover image components that are unobservable in the data [22, Chapter 3.6].

More interesting is the case when F is chosen as a gradient or Laplacian operator, because $\|Fx\|_2^2$ provides a measure of the variability of x. So doing, the regularity constraint forces solutions with limited high-frequency energy, and thus it captures the prior belief that the recovered image should be smooth. Interestingly, it can be shown that such a regularization leads to solutions containing image components that are unobservable in the data [22, Chapter 3.6].

Regardless the choice of F, the use of a quadratic regularization results in a linear filtering of the observed data, as there exists a Lagrangian equivalence between Problem (3.2) and the regularized least-squares method in (1.3). Consequently, when F is chosen so as to suppress the effect of high-frequency noise, the linear filtering also reduces the high-frequency energy of the true image, leading to an undesirable smoothing of object edges in the recovered image.

3.1.2 Total variation regularization

A powerful approach to overcome the limitation of the quadratic regularization consists of resorting to the *total variation* penalty [61, 195], whose discretized version is defined as the $\boldsymbol{\ell}_{1,p}$-norm (with $p \geq 0$) of the image gradient coefficients

When ℓ is located in the last column of the image, $n_{\ell,1} = \ell$ so that $x^{(\ell)} - x^{(n_{\ell,1})} = 0$. Likewise, $n_{\ell,2} = \ell$ on the last row.

$$\mathrm{TV}_p(x) = \sum_{\ell=1}^{N} \left(|x^{(\ell)} - x^{(n_{\ell,1})}|^p + |x^{(\ell)} - x^{(n_{\ell,2})}|^p \right)^{1/p}, \tag{3.4}$$

where $(n_{\ell,1}, n_{\ell,2}) \in \{1, \ldots, N\}^2$ denote the positions of the horizontal/vertical nearest neighbors of $x^{(\ell)}$. The above function actually computes the total amount of change that the image x goes through, and thus it provides a robust measure of the variability of x. As a result, this regularization tends to prefer solutions with localized steep gradients, leading to the suppression of high-frequency details in the recovered image while preserving the main geometry of the true image.

The TV regularization can be plugged into Problem (3.2) by setting the transformation F as the 2D gradient operator, namely

$$Fx = \begin{bmatrix} x^{(1)} - x^{(n_{1,1})} \\ x^{(1)} - x^{(n_{1,2})} \\ \vdots \\ x^{(N)} - x^{(n_{N,1})} \\ x^{(N)} - x^{(n_{N,2})} \end{bmatrix} \begin{matrix} \left.\vphantom{\begin{matrix} a \\ b \end{matrix}}\right\} \mathsf{y}^{(1)} \in \mathbb{R}^2 \\ \vdots \\ \left.\vphantom{\begin{matrix} a \\ b \end{matrix}}\right\} \mathsf{y}^{(N)} \in \mathbb{R}^2, \end{matrix} \tag{3.5}$$

and the function h in (3.3) with $M = 2N$, $q = 1$, and $L = N$. The choices $p = 1$ or $p = 2$ are commonly referred to as *anisotropic* or *isotropic* TV, respectively.

3.1.3 Nonlocal total variation regularization

Total variation has emerged as a successful mixed-norm regularization, with $p \in \{0, 1, 2, +\infty\}$ being the most-used $\boldsymbol{\ell}_{1,p}$-norms [28, 171, 227, 228]. However, TV fails to preserve textures, details and fine structures, because they are hardly distinguishable from high-frequency noise. To improve this behavior, the TV model has been extended by using some generalizations based on higher-order spatial differences [27, 177], higher-degree directional derivatives [118, 119], or the non-locality principle [105, 205, 229]. Another approach to overcome this limitation is to replace the gradient with an operator yielding a better sparse representation of the image, such as a frame [39, 124, 164] or learned dictionary [3, 92, 122, 163]. In this context, the family of BM3D algorithms [79, 162] has been recently formulated in terms of analysis and synthesis frames [80], substantiating the use of non-locality as a valuable image modeling tool.

The *nonlocal total variation* extends the TV model by measuring the interactions of each pixel with a "selected" subset of pixels chosen in the whole image, yielding

$$\mathrm{NLTV}_{q,p}(x) = \sum_{\ell=1}^{N} \left(\sum_{n \in \mathcal{N}_\ell} \omega_{\ell,n}^p \, |x^{(\ell)} - x^{(n)}|^p \right)^{q/p}, \tag{3.6}$$

where, for every $\ell \in \{1, \ldots, N\}$, the set $\mathcal{N}_\ell \subset \{1, \ldots, N\} \setminus \{\ell\}$ contains some positions located around ℓ, and $(\omega_{\ell,n})_{n \in \mathcal{N}_\ell}$ are positive weights that measure the *similarity* between $x^{(\ell)}$ and its neighbors $(x^{(n)})_{n \in \mathcal{N}_\ell}$. Several types of $\boldsymbol{\ell}_{q,p}$-norms can be used to penalize the nonlocal variations. The case $q = p = 2$ may be seen as a variational extension of the *nonlocal-means* algorithm [32], while the choice $q = p = 1$ leads to the *nonlocal-medians* algorithm [150]. Finally, $q = 1$ and $p \in \{2, +\infty\}$ yields the common usage as mixed-norm regularization.

The NLTV regularization can be plugged into Problem (3.2) by setting the transformation F as the nonlocal gradient operator, namely

$$Fx = \begin{bmatrix} \big[\omega_{1,n}(x^{(1)} - x^{(n)})\big]_{n \in \mathcal{N}_1} \\ \vdots \\ \big[\omega_{N,n}(x^{(N)} - x^{(n)})\big]_{n \in \mathcal{N}_N} \end{bmatrix} \begin{matrix} \} \, \mathsf{y}^{(1)} \in \mathbb{R}^{M_1} \\ \vdots \\ \} \, \mathsf{y}^{(N)} \in \mathbb{R}^{M_N}, \end{matrix} \tag{3.7}$$

The nonlocal operator reduces to the classical gradient operator when $\mathcal{N}_\ell$ only contains the horizontal/vertical nearest neighbors of ℓ and $\omega_{\ell,n} \equiv 1$.

and the function h in (3.3) with $M = \sum_{\ell=1}^{N} M_\ell$ and $L = N$. For every $(\ell, n) \in \{1, \ldots, N\}^2$, the weight $\omega_{\ell,n} > 0$ depends on the similarity between patches built around the pixels ℓ and n of the true image. Since the degradation process in (3.1) may involve some missing data, a two-step approach must be undertaken in order to estimate these weights. In the first step, the TV regularization is used in order to obtain an estimate $\widetilde{x}$ of the target image. This estimate serves in the second step to compute the weights through a *self-similarity* measure:

$$\omega_{\ell,n} = \widetilde{\omega}_\ell \exp\left(-\delta^{-2} \, \|\mathrm{L}_\ell \widetilde{\mathrm{x}} - \mathrm{L}_n \widetilde{\mathrm{x}}\|_2^2\right), \tag{3.8}$$

where $\delta > 0$, L_ℓ (resp. L_n) selects a $\widetilde{Q} \times \widetilde{Q}$ patch centered at position ℓ (resp. n) after a linear processing depending on the position ℓ (resp. n) [32, 99], and the constant $\widetilde{\omega}_\ell > 0$ is set so as to normalize the weights (i.e., $\sum_{n \in \mathcal{N}_\ell} \omega_{\ell,n} = 1$). For every $\ell \in \{1, \ldots, N\}$, the neighborhood $\mathcal{N}_\ell$ is built according to the procedure described in [104], which allows one to limit the size M_ℓ of the neighborhood to a fixed value $\overline{M}$ (in all the experiments, we set $\widetilde{Q} = 5$, $\delta = 35$ and $\overline{M} = 14$).

3.2 Optimization method

The standard way of solving Problem (3.2) via proximal methods leads to a sequence of steps involving the projection onto the $\boldsymbol{\ell}_{q,p}$-ball, as illustrated in Algorithms 3.1 and 3.3 with M+LFBF [66] and SDMM [65]. However, an expression of this projection is not available in the cases relevant to image recovery, that are $q = 1$ and $p \in \{1, 2, +\infty\}$. Hence, it must be computed through specific numerical method [71, 188, 217], as explained in Section 1.3.3.

An alternative approach that avoids the use of inner iterations consists of reformulating Problem (3.2) by the proposed epigraphical splitting, yielding

$$\underset{(x,\zeta)\in[0,255]^N\times V}{\text{minimize}} \quad \|Ax - z\|_2^2 \quad \text{s.t.} \quad (Fx, \zeta) \in E, \tag{3.9}$$

where the convex sets E and V are defined in Sections 3.2.1 and 3.2.2 according to the type of $\boldsymbol{\ell}_{q,p}$-ball involved in Problem (3.2). The iterations of M+LFBF and SDMM associated to the above problem are reported in Algorithms 3.2 and 3.4, which involve only projections available in closed form, as explained next.

3.2.1 Mixed norms

Let $q = 1$ and $p \in \{2, +\infty\}$ in (3.3). Then, the sublevel set of the $\boldsymbol{\ell}_{1,p}$-norm can be decomposed by introducing an auxiliary vector $\zeta = (\zeta^{(\ell)})_{1\leq\ell\leq L}$ such that

$$\sum_{\ell=1}^{L} \|\mathsf{y}^{(\ell)}\|_p \leq \eta \quad \Leftrightarrow \quad \begin{cases} (\forall \ell \in \{1, \dots, L\}) & \|\mathsf{y}^{(\ell)}\|_p \leq \zeta^{(\ell)}, \\ \displaystyle\sum_{\ell=1}^{L} \zeta^{(\ell)} \leq \eta. \end{cases} \tag{3.10}$$

Consequently, the sets E and V in Problem (3.9) are defined as follows:

$$\begin{aligned} E &= \big\{(y, \zeta) \in \mathbb{R}^M \times \mathbb{R}^L \;\big|\; (\forall \ell \in \{1, \dots, L\}) \quad (\mathsf{y}^{(\ell)}, \zeta^{(\ell)}) \in \operatorname{epi} \|\cdot\|_p\big\}, \\ V &= \big\{\zeta \in \mathbb{R}^L \;\big|\; 1_L^\top \zeta \leq \eta\big\}, \end{aligned} \tag{3.11}$$

with the epigraphical projections given in Corollary 2.3.5 and Proposition 2.3.7.

3.2.2 Usual norms

Let $q \geq 1$ and $p = q$ in (3.3). Then, the sublevel set of the $\boldsymbol{\ell}_{q,q}$-norm can be split by introducing an auxiliary vector $\zeta = (\zeta^{(\ell,m)})_{1\leq\ell\leq L, 1\leq m\leq M_\ell}$ such that

$$\sum_{\ell=1}^{L} \|\mathsf{y}^{(\ell)}\|_q^q \leq \eta \;\; \Leftrightarrow \;\; \begin{cases} (\forall \ell \in \{1, \dots, L\})(\forall m \in \{1, \dots, M_\ell\}) & |y^{(\ell,m)}|^q \leq \zeta^{(\ell,m)}, \\ \displaystyle\sum_{\ell=1}^{L} \sum_{m=1}^{M_\ell} \zeta^{(\ell,m)} \leq \eta. \end{cases} \tag{3.12}$$

Consequently, the sets E and V in Problem (3.9) are defined as follows:

$$\begin{aligned} E = \{(y, \zeta) \in \mathbb{R}^M \times \mathbb{R}^M \;\mid\; & (\forall \ell \in \{1, \dots, L\}) \\ & (\forall m \in \{1, \dots, M_\ell\}) \quad (y^{(\ell,m)}, \zeta^{(\ell,m)}) \in \operatorname{epi} |\cdot|^q\big\}, \end{aligned} \tag{3.13}$$

$$V = \big\{\zeta \in \mathbb{R}^M \;\big|\; 1_M^\top \zeta \leq \eta\big\}, \tag{3.14}$$

and the epigraphical projections are given in Propositions 2.3.2 and 2.3.3.

Algorithm 3.1 *M+LFBF algorithm [66] specialized to Problem* (3.2)

INITIALIZATION

$$\left\lfloor \begin{array}{l} \text{choose } \left(x^{[0]}, y^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^M \\ \text{set } \gamma \in \left]0, \left(2\|A\|^2 + \|F\|^2\right)^{-1/2}\right[\end{array} \right.$$

FOR $i = 0, 1, \ldots$

$$\left\lfloor \begin{array}{l} p^{[i]} = P_{[0,255]^N}\left(x^{[i]} - \gamma\left(2A^\top A x^{[i]} - 2A^\top z + F^\top y^{[i]}\right)\right) \\ \widehat{y}^{[i]} = y^{[i]} + \gamma F x^{[i]} \\ v^{[i]} = \widehat{y}^{[i]} - \gamma P_{\mathrm{lev}_{\leq \eta}\, h}\left(\widehat{y}^{[i]}/\gamma\right) \\ y^{[i+1]} = v^{[i]} + \gamma F\left(p^{[i]} - x^{[i]}\right) \\ x^{[i+1]} = p^{[i]} - \gamma\Big(2A^\top A(p^{[i]} - x^{[i]}) + F^\top(v^{[i]} - y^{[i]})\Big) \end{array} \right.$$

Algorithm 3.2 *M+LFBF algorithm [66] specialized to Problem* (3.9)

INITIALIZATION

$$\left\lfloor \begin{array}{l} \text{choose } \left(x^{[0]}, \zeta^{[0]}\right) \in \mathbb{R}^N \times V \\ \text{choose } \left(y^{[0]}, \xi^{[0]}\right) \in \mathbb{R}^M \times V \\ \text{set } \gamma \in \left]0, \left(2\|A\|^2 + \max\{\|F\|^2, 1\}\right)^{-1/2}\right[\end{array} \right.$$

FOR $i = 0, 1, \ldots$

$$\left\lfloor \begin{array}{l} p^{[i]} = P_{[0,255]^N}\left(x^{[i]} - \gamma\left(2A^\top A x^{[i]} - 2A^\top z + F^\top y^{[i]}\right)\right) \\ \rho^{[i]} = P_V\left(\zeta^{[i]} - \gamma\,\xi^{[i]}\right) \\ (\widehat{y}^{[i]}, \widehat{\xi}^{[i]}) = (y^{[i]}, \xi^{[i]}) + \gamma\,(Fx^{[i]}, \zeta^{[i]}) \\ (v^{[i]}, \nu^{[i]}) = (\widehat{y}^{[i]}, \widehat{\xi}^{[i]}) - \gamma P_E\left(\widehat{y}^{[i]}/\gamma, \widehat{\xi}^{[i]}/\gamma\right) \\ y^{[i+1]} = v^{[i]} + \gamma F\left(p^{[i]} - x^{[i]}\right) \\ \xi^{[i+1]} = \nu^{[i]} + \gamma\left(\rho^{[i]} - \zeta^{[i]}\right) \\ x^{[i+1]} = p^{[i]} - \gamma\Big(2A^\top A(p^{[i]} - x^{[i]}) + F^\top(v^{[i]} - y^{[i]})\Big) \\ \zeta^{[i+1]} = \rho^{[i]} - \gamma\left(\nu^{[i]} - \xi^{[i]}\right) \end{array} \right.$$

Algorithm 3.3 *SDMM algorithm [65] specialized to Problem* (3.2)

INITIALIZATION

$$
\left\lfloor
\begin{array}{l}
\text{choose } \left(y_1^{[0]}, y_2^{[0]}, y_3^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^K \times \mathbb{R}^M \\
\text{choose } \left(\overline{y}_1^{[0]}, \overline{y}_2^{[0]}, \overline{y}_3^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^K \times \mathbb{R}^M \\
\text{set } \gamma \in]0, +\infty[\\
\text{define } H = \mathrm{Id} + A^\top A + F^\top F
\end{array}
\right.
$$

FOR $i = 0, 1, \ldots$

$$
\left\lfloor
\begin{array}{l}
x^{[i]} = H^{-1}\left(y_1^{[i]} - \overline{y}_1^{[i]} + A^\top(y_2^{[i]} - \overline{y}_2^{[i]}) + F^\top(y_3^{[i]} - \overline{y}_3^{[i]})\right) \\
y_1^{[i+1]} = P_{[0,255]^N}\left(x^{[i]} + \overline{y}_1^{[i]}\right) \\
y_2^{[i+1]} = \mathrm{prox}_{\gamma\|\cdot - z\|_2^2}\left(Ax^{[i]} + \overline{y}_2^{[i]}\right) \\
y_3^{[i+1]} = P_{\mathrm{lev}_{\leq \eta}\, h}\left(Fx^{[i]} + \overline{y}_3^{[i]}\right) \\
\overline{y}_1^{[i+1]} = \overline{y}_1^{[i]} + \quad x^{[i]} - y_1^{[i+1]} \\
\overline{y}_2^{[i+1]} = \overline{y}_2^{[i]} + Ax^{[i]} - y_2^{[i+1]} \\
\overline{y}_3^{[i+1]} = \overline{y}_3^{[i]} + Fx^{[i]} - y_3^{[i+1]}
\end{array}
\right.
$$

Algorithm 3.4 *SDMM algorithm [65] specialized to Problem* (3.9)

INITIALIZATION

$$
\left\lfloor
\begin{array}{l}
\text{choose } \left(y_1^{[0]}, y_2^{[0]}, y_3^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^K \times \mathbb{R}^M \\
\text{choose } \left(\overline{y}_1^{[0]}, \overline{y}_2^{[0]}, \overline{y}_3^{[0]}\right) \in \mathbb{R}^N \times \mathbb{R}^K \times \mathbb{R}^M \\
\text{choose } \left(\nu_2^{[0]}, \overline{\nu}_2^{[0]}, \nu_3^{[0]}, \overline{\nu}_3^{[0]}\right) \in V \times V \times V \times V \\
\text{set } \gamma \in]0, +\infty[\\
\text{define } H = \mathrm{Id} + A^\top A + F^\top F
\end{array}
\right.
$$

FOR $i = 0, 1, \ldots$

$$
\left\lfloor
\begin{array}{ll}
x^{[i]} & = H^{-1}\left(y_1^{[i]} - \overline{y}_1^{[i]} + A^\top(y_2^{[i]} - \overline{y}_2^{[i]}) + F^\top(y_3^{[i]} - \overline{y}_3^{[i]})\right) \\
\zeta^{[i]} & = \nu_2^{[i]} - \overline{\nu}_2^{[i]} + \nu_3^{[i]} - \overline{\nu}_3^{[i]} \\
y_1^{[i+1]} & = P_{[0,255]^N}\left(x^{[i]} + \overline{y}_1^{[i]}\right) \\
y_2^{[i+1]} & = \mathrm{prox}_{\gamma\|\cdot - z\|_2^2}\left(Ax^{[i]} + \overline{y}_2^{[i]}\right) \\
\nu_2^{[i+1]} & = P_V\left(\zeta^{[i]} + \overline{\nu}_2^{[i]}\right) \\
(y_3^{[i+1]}, \nu_3^{[i+1]}) & = P_E\left(Fx^{[i]} + \overline{y}_3^{[i]}, \zeta^{[i]} + \overline{\nu}_3^{[i]}\right) \\
\overline{y}_1^{[i+1]} & = \overline{y}_1^{[i]} + x^{[i]} - y_1^{[i+1]} \\
(\overline{y}_2^{[i+1]}, \overline{\nu}_2^{[i+1]}) & = (\overline{y}_2^{[i]}, \overline{\nu}_2^{[i]}) + (Ax^{[i]}, \zeta^{[i]}) - (y_2^{[i+1]}, \nu_2^{[i+1]}) \\
(\overline{y}_3^{[i+1]}, \overline{\nu}_3^{[i+1]}) & = (\overline{y}_3^{[i]}, \overline{\nu}_3^{[i]}) + (Fx^{[i]}, \zeta^{[i]}) - (y_3^{[i+1]}, \nu_3^{[i+1]})
\end{array}
\right.
$$

3.3 Numerical results

A numerical analysis of the proposed approach is now conducted. In Section 3.3.1, the visual impact of TV_p and $\mathrm{NLTV}_{1,p}$ regularization is evaluated for different choices of $p \in \{1, 2, +\infty\}$. In Section 3.3.2, the execution times attained with the epigraphical approach illustrated in Algorithms 3.2 and 3.4 are compared to those achieved with the standard approach given in Algorithms 3.1 and 3.3.

3.3.1 Assessment of the recovered quality

Figure 3.1 shows two images degraded with the model in (3.1). The same images recovered by solving Problem (3.2) are illustrated in Figures 3.2 and 3.3 for several types of regularity constraints, in which the corresponding bounds η were hand-tuned in order to achieve the best SNR values. Additional examples with various grayscale images are reported in Table 3.1. The results show the interest of considering nonlocal operators for modeling the regularity present in natural images. They also confirm that $\mathrm{NLTV}_{1,2}$ achieves a better performance with grayscale images, whereas $\mathrm{NLTV}_{1,\infty}$ is better suited for color images [171]. Note however that the regularization is applied separately on each color channel. In this regard, a more powerful approach will be proposed in Chapter 5.

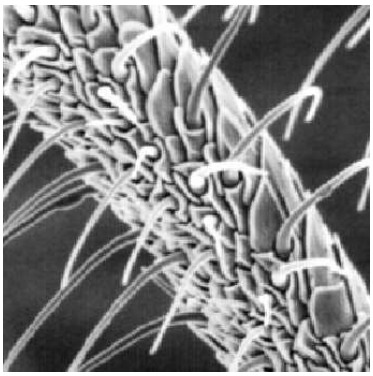

(a) Original image ($N = 256 \times 256$).

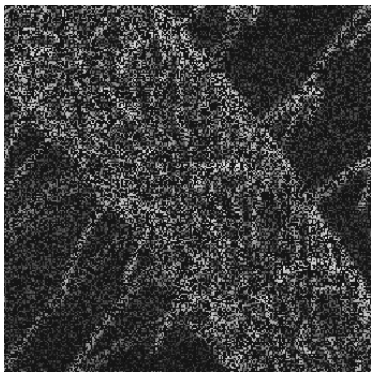

(b) Degraded image ($K = 0.4 \times N$).

(c) Original image ($N = 320 \times 480 \times 3$).

(d) Degraded image ($K = 0.4 \times N$).

Figure 3.1 *Examples of images degraded by the model* (3.1) *through the use of a* 3×3 *uniform blur, a decimation that randomly removes* 60% *of the pixels, and an additive white Gaussian noise with variance* 10^2. *For color images, each channel is degraded separately.*

Table 3.1 *SNR (dB) and SSIM indexes (blur:* 3×3*, noise:* $\sigma = 10$*, decimation:* 60%*)*

SNR (dB) – SSIM	N	TV_2	TV_∞	$\mathrm{NLTV}_{1,2}$	$\mathrm{NLTV}_{1,\infty}$
Lena	256^2	23.18 – 0.783	22.77 – 0.769	**24.18 – 0.812**	24.14 – 0.812
Boat	256^2	20.25 – 0.739	19.74 – 0.718	**21.13 – 0.770**	20.77 – 0.741
Cameraman	256^2	20.06 – 0.774	19.68 – 0.755	**20.71 – 0.801**	20.17 – 0.743
House	256^2	25.47 – 0.823	24.70 – 0.808	**26.31 – 0.836**	25.87 – 0.823
Man	256^2	19.24 – 0.725	18.96 – 0.714	**19.66 – 0.741**	19.51 – 0.736
Peppers	512^2	23.69 – 0.801	23.25 – 0.786	**24.80 – 0.829**	24.45 – 0.813
Barbara	512^2	16.74 – 0.653	16.64 – 0.642	**17.02 – 0.673**	16.99 – 0.652
Hill	512^2	22.18 – 0.723	21.89 – 0.715	**22.55 – 0.735**	22.43 – 0.733

3.3.2 Assessment of the execution time

Figure 3.4 plots the relative distance $\|x^{[i]} - x^{[\infty]}\|/\|x^{[\infty]}\|$ (as a function of time) of iterates generated by Algorithm 3.1 (*M+LFBF direct*), Algorithm 3.2 (*M+LFBF epi*), Algorithm 3.3 (*SDMM direct*), and Algorithm 3.4 (*SDMM epi*). This experiment refers to the example in Figure 3.3, where the stopping criterion was set to $\|x^{[i+1]} - x^{[i]}\| \leq 10^{-4}\|x^{[i]}\|$, with $x^{[\infty]}$ denoting the solution obtained when such a criterion is reached up to 10^6 iterations (note that $x^{[\infty]}$ may not be unique, hence it was computed for each algorithm independently). These plots indicate that the proposed epigraphical approach yields a faster execution time.

Since the constraint bound η may not be known precisely, it is important to evaluate the impact of its choice on the performance of the proposed epigraphical splitting (it is outside the scope of this thesis to devise an optimal strategy to set this bound). In Tables 3.2-3.5, we compare the execution times of epigraphical and direct approaches for different choices of regularization constraints and values of η. This experiment refers to the example in Figure 3.2, as the computational burden of the projection onto the $\boldsymbol{\ell}_{1,\infty}$-ball becomes unbearable with an image bigger than 256×256 (see Figure 3.4). The stopping criterion is set to $\|x^{[i+1]} - x^{[i]}\| \leq 10^{-4}\|x^{[i]}\|$. For more readability, the values of η are given as a multiplicative factor of TV_p or $\mathrm{NLTV}_{1,p}$ evaluated on the original image.

- Tables 3.2 and 3.3 report the comparisons for TV_2 and TV_∞, respectively. The execution times indicate that the epigraphical approach yields a faster convergence although it requires more iterations in order to converge. This can be explained by the computational cost of the subiterations required by the direct projections onto the $\boldsymbol{\ell}_{1,p}$-ball.

- Tables 3.2 and 3.3 also show that errors within $\pm 20\%$ from the optimal value for η lead to SNR variations within 2%.

- Tables 3.4 and 3.5 collect the results of $\mathrm{NLTV}_{1,2}$ and $\mathrm{NLTV}_{1,\infty}$ for different neighborhood sizes. The execution times show that the epigraphical approach is faster than the direct one for both considered algorithms.

- Tables 3.4 and 3.5 also show that errors within $\pm 20\%$ from the optimal bound value lead to SNR variations within 1%.

The codes used in all the experiments were developed in MATLAB R2011b (with the operators F and $F^\top$ being implemented in C throuh mex-files) and executed on an Intel Xeon CPU at 2.80 GHz and 8 GB of RAM.

(a) Original image (zoom).

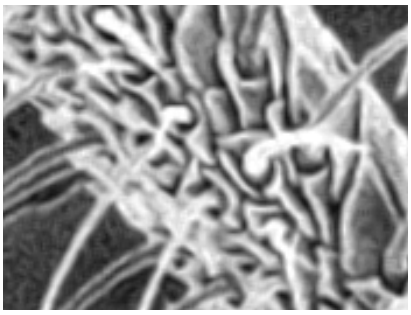

(b) $\mathrm{NLTV}_{2,2}$: 22.26–0.842.

(c) TV_1: 19.79–0.838.

(d) TV_2: 20.80–0.855.

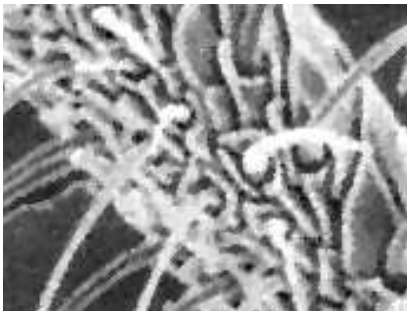

(e) TV_∞: 20.25–0.853.

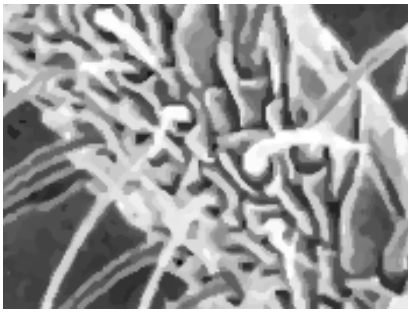

(f) $\mathrm{NLTV}_{1,1}$: 20.93–0.865.

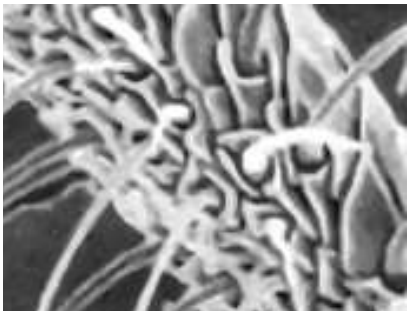

(g) $\mathrm{NLTV}_{1,2}$: **22.62–0.897**.

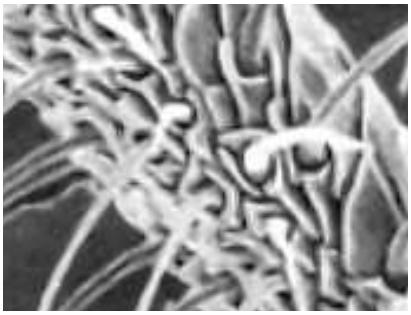

(h) $\mathrm{NLTV}_{1,\infty}$: 22.38–0.897.

Figure 3.2 *SNR (dB) – SSIM indexes for the recovery of a grayscale image degraded with a uniform blur of size* 3×3*, a noise of variance* 10^2*, and* 60% *of decimation.*

(a) Original image (zoom).

(b) $\mathrm{NLTV}_{2,2}$: 18.12–0.761.

(c) TV_1: 17.78–0.787.

(d) TV_2: 18.36–0.821.

(e) TV_∞: 18.91–0.824.

(f) $\mathrm{NLTV}_{1,1}$: 18.93–0.828.

(g) $\mathrm{NLTV}_{1,2}$: 19.47–0.839.

(h) $\mathrm{NLTV}_{1,\infty}$: **20.17–0.847**.

Figure 3.3 *SNR (dB) – SSIM indexes for the recovery of a color image degraded with a uniform blur of size* 3×3*, a noise of variance* 10^2*, and* 60% *of decimation.*

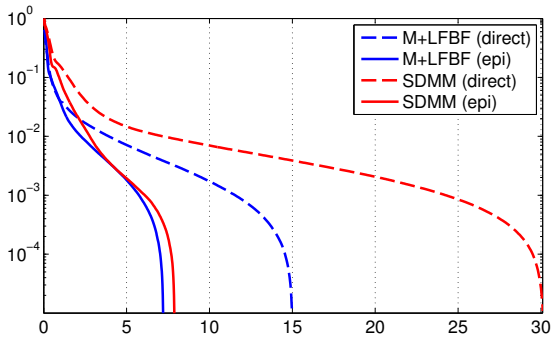

(a) TV_2.

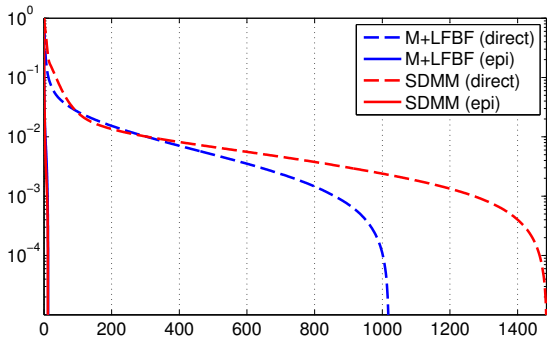

(b) TV_∞.

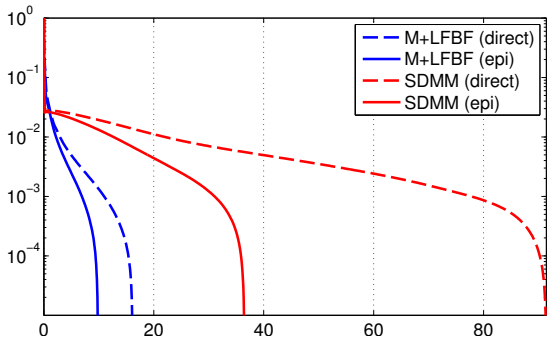

(c) $\mathrm{NLTV}_{1,2}$.

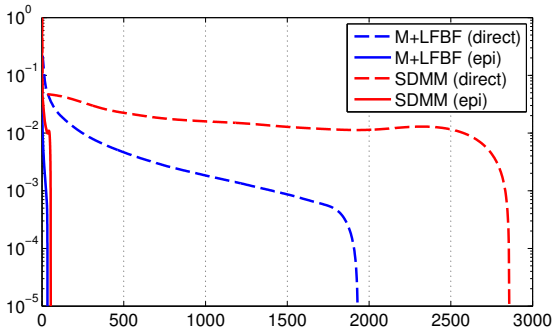

(d) $\mathrm{NLTV}_{1,\infty}$.

Figure 3.4 *Relative distance to $x^{[\infty]}$ vs execution time (seconds): plots comparing the epigraphical and direct approaches implemented with M+LFBF and SDMM.*

Table 3.2 *Results for the TV_2 constraint and different values of η.*

η	SNR (dB) – SSIM	SDMM					M+LFBF				
		direct		epigraphical		speed up	direct		epigraphical		speed up
		# iter.	sec.	# iter.	sec.		# iter.	sec.	# iter.	sec.	
0.45	19.90 – 0.733	107	6.07	174	2.03	2.99	113	6.15	182	3.49	1.76
0.50	20.18 – 0.745	117	6.95	159	1.95	3.57	116	6.97	168	3.44	2.03
0.56	**20.23 – 0.745**	**129**	**8.36**	**153**	**1.90**	**4.41**	**124**	**8.17**	**159**	**3.01**	**2.72**
0.62	20.16 – 0.737	141	9.44	155	1.83	5.16	131	8.62	159	3.26	2.65
0.67	20.00 – 0.724	154	10.20	162	2.17	4.71	140	10.00	164	2.84	3.52

Table 3.3 *Results for the TV_∞ constraint and different values of η.*

η	SNR (dB) – SSIM	SDMM					M+LFBF				
		direct		epigraphical		speed up	direct		epigraphical		speed up
		# iter.	sec.	# iter.	sec.		# iter.	sec.	# iter.	sec.	
0.45	19.52 – 0.726	160	312.55	231	3.89	80.43	183	347.10	252	6.43	53.96
0.50	19.71 – 0.728	168	342.01	215	3.75	91.31	185	368.24	236	5.83	63.17
0.56	**19.71 – 0.734**	**180**	**373.60**	**211**	**3.49**	**106.93**	**189**	**386.29**	**229**	**5.53**	**69.91**
0.62	19.59 – 0.715	196	412.68	216	3.67	112.50	198	411.04	229	5.86	70.15
0.67	19.39 – 0.698	211	448.77	223	3.76	119.27	207	437.66	234	5.76	75.96

Table 3.4 *Results obtained with $NLTV_{1,2}$ constraint and different values of η*

η	SNR (dB) – SSIM	SDMM					M+LFBF				
		direct		epigraphical		speed up	direct		epigraphical		speed up
		# iter.	sec.	# iter.	sec.		# iter.	sec.	# iter.	sec.	
Neighbourhood size: 3×3											
0.43	20.82 – 0.757	208	20.67	211	10.93	1.89	82	6.95	93	3.76	1.85
0.49	20.97 – 0.765	167	16.84	177	9.01	1.87	75	6.61	83	3.47	1.91
0.54	21.02 – 0.767	147	15.31	157	7.93	1.93	71	6.45	77	3.15	2.04
0.59	20.98 – 0.764	134	14.44	148	7.67	1.88	72	6.58	77	3.24	2.03
0.65	20.88 – 0.757	133	14.82	136	7.11	2.08	76	7.53	80	3.27	2.30
Neighbourhood size: 5×5											
0.43	21.00 – 0.766	301	56.03	343	45.18	1.24	82	8.51	90	5.43	1.57
0.49	21.15 – 0.773	260	49.03	302	39.64	1.24	75	7.90	81	4.90	1.61
0.54	**21.20 – 0.775**	**242**	**46.31**	**283**	**37.72**	**1.23**	**71**	**8.26**	**75**	**4.47**	**1.85**
0.59	21.17 – 0.773	231	46.20	268	36.56	1.26	70	7.94	74	4.49	1.77
0.65	21.08 – 0.767	220	44.64	252	34.46	1.30	73	8.40	76	4.59	1.83

Table 3.5 *Results obtained with $NLTV_{1,\infty}$ and different values of η*

η	SNR (dB) – SSIM	SDMM					M+LFBF				
		direct		epigraphical		speed up	direct		epigraphical		speed up
		# iter.	sec.	# iter.	sec.		# iter.	sec.	# iter.	sec.	
Neighbourhood size: 3×3											
0.43	20.78 – 0.762	434	1470.46	449	25.03	58.76	225	730.26	244	12.35	59.15
0.49	20.86 – 0.764	395	1319.64	413	22.86	57.72	221	692.25	237	11.92	58.08
0.54	20.83 – 0.760	363	1193.61	382	21.46	55.62	217	667.50	233	11.46	58.22
0.59	20.73 – 0.752	340	1093.26	354	19.77	55.30	216	653.79	230	11.67	56.01
0.65	20.58 – 0.740	322	1007.55	336	18.64	54.06	216	643.00	229	11.45	56.18
Neighbourhood size: 5×5											
0.43	20.91 – 0.769	384	2069.62	452	64.42	32.13	233	863.01	252	18.47	46.73
0.49	20.97 – 0.767	326	1700.34	412	58.66	28.99	231	822.06	247	18.36	44.77
0.54	**20.98 – 0.771**	**290**	**1476.98**	**389**	**55.35**	**26.69**	**229**	**787.61**	**245**	**17.90**	**43.99**
0.59	20.88 – 0.759	276	1336.16	374	52.64	25.38	230	772.42	245	17.57	43.96
0.65	20.75 – 0.749	268	1220.14	362	51.45	23.72	231	760.86	245	17.81	42.72

3.4 Conclusions

In this chapter, we have applied the epigraphical splitting to image recovery problems involving a regularity constraint based on mixed norms. On the one hand, the obtained results demonstrate the good performance of nonlocal measures and the interest of considering the infinity norm for the regularization of color images. On the other hand, the obtained results indicate that the epigraphical splitting leads to execution times faster than those achieved with the direct computation of the projections via standard iterative solutions. Parallelization of our codes should even allow us to accelerate them [102].

In the next chapter, we will employ the epigraphical splitting to develop approximation methods for addressing more general convex constraints.

Le savant n'est pas l'homme qui fournit les vraies réponses: c'est celui qui pose les vraies questions.

Claude Lévi-Strauss

Chapter 4
Constraints based on piecewise-affine functions

The main focus of this chapter is the epigraphical splitting applied to outer-approximated constraints. To this end, we present a primal-dual approach grounded on the epigraphical projection of the max function. Experiments with constraints based on Kullback-Leibler divergence and logistic loss demonstrate the efficiency of our approach.

4.1 Introduction

The epigraphical splitting revolves around the idea of replacing the sublevel set of a block-separable function $h \in \Gamma_0(\mathbb{R}^M)$ with a collection of epigraphs, namely

$$h(y) = \sum_{\ell=1}^{L} h_\ell(\mathrm{y}^{(\ell)}) \leq \eta \qquad \Leftrightarrow \qquad \begin{cases} 1_L^\top \zeta \leq \eta, \\ (\forall \ell \in \{1, \ldots, L\}) \quad h_\ell(\mathrm{y}^{(\ell)}) \leq \zeta^{(\ell)}, \end{cases} \tag{4.1}$$

with the vector $y \in \mathbb{R}^M$ block-decomposed as in (2.3). As illustrated in Chapter 3, the above decomposition turns out to be very efficient when the functions $(h_\ell)_{1 \leq \ell \leq L}$ are $\boldsymbol{\ell}_p$-norms with $p \in \{1, 2, +\infty\}$, since the associated epigraphical projections can be quickly evaluated. However, the same efficiency cannot be claimed when there is no available expression for the epigraphical projection.

The present chapter proposes a solution to overcome the aforementioned difficulty. The principle is to outer approximate the sublevel set of h by replacing, for every $\ell \in \{1, \ldots, L\}$, the term h_ℓ with a lower-approximating function having a tractable epigraph, so as to facilitate the use of the epigraphical splitting. This can be certainly achieved by means of piecewise-affine approximations, since the associated epigraph can be handled through the epigraphical projection of the max function derived in Section 2.3.3. Indeed, a piecewise-affine function can be expressed as the maximum of a number of scalar products, namely

$$(\forall \mathrm{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad \widehat{h}_\ell(\mathrm{y}^{(\ell)}) = \max_{1 \leq j \leq J_\ell} \delta_{\ell,j}^\top \, \mathrm{y}^{(\ell)} + \mu_{\ell,j}, \tag{4.2}$$

where $\delta_{\ell,j} \in \mathbb{R}^{M_\ell}$ and $\mu_{\ell,j} \in \mathbb{R}$ for every $j \in \{1, \ldots, J_\ell\}$, with $J_\ell \in \mathbb{N}^*$.

The chapter is organized as follows. Section 4.2 details an approximation technique based on piecewise-affine functions, together with some theoretical equivalences between the original problem and the approximated one. Since the epigraphical splitting decomposes the outer-approximated constraint into a collection of convex polyhedrons, Section 4.3 presents a primal-dual approach to deal with such constraints through the epigraphical projection of the max function. Finally, Section 4.4 demonstrates the validity of the proposed approximation method through an image recovery problem involving the Kullback-Leibler divergence, and a supervised classification problem based on logistic regression.

4.2 Lower approximation

The proposed approximation deals with optimization problems of the form

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\,x) \quad \text{s.t.} \quad h(Fx) = \sum_{\ell=1}^{L} h_\ell\,(F_\ell\,x) \leq \eta, \tag{4.3}$$

and it consists of replacing each term h_ℓ with a function in $\Gamma_0(\mathbb{R}^{M_\ell})$ such that

$$(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad \widehat{h}_\ell(\mathsf{y}^{(\ell)}) \leq h_\ell(\mathsf{y}^{(\ell)}). \tag{4.4}$$

So doing, the constraint in Problem (4.3) can be outer approximated as follows

$$\big\{y \in \mathbb{R}^M \;\big|\; \sum_{\ell=1}^{L} h_\ell(\mathsf{y}^{(\ell)}) \leq \eta\big\} \quad \subset \quad \big\{y \in \mathbb{R}^M \;\big|\; \sum_{\ell=1}^{L} \widehat{h}_\ell(\mathsf{y}^{(\ell)}) \leq \eta\big\}, \tag{4.5}$$

which eventually leads to the approximated problem

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\,x) \quad \text{s.t.} \quad \sum_{\ell=1}^{L} \widehat{h}_\ell\,(F_\ell\,x) \leq \eta. \tag{4.6}$$

A possible choice for the function $\widehat{h}_\ell$ is illustrated in Section 4.2.1, whereas a theoretical equivalence between the above problems is detailed in Section 4.2.2.

4.2.1 Piecewise-affine function

In the spirit of cutting plane methods [133] and some of their extensions [57, 116], the function h_ℓ in (4.3) can be lower approximated by a piecewise-affine function, as shown in Figure 4.1. To do so, one can notice that by conjugation

$$h_\ell(\mathsf{y}^{(\ell)}) = \sup_{\delta_\ell\in\mathbb{R}^{M_\ell}} \delta_\ell^\top \mathsf{y}^{(\ell)} - h_\ell^*(\delta_\ell). \tag{4.7}$$

By computing the supremum over a number J_ℓ of vectors $\{\delta_{\ell,1},\dots,\delta_{\ell,J_\ell}\} \subset \mathbb{R}^{M_\ell}$, one can lower approximate h_ℓ through a function of the form

$$\widehat{h}_\ell(\mathsf{y}^{(\ell)}) = \max_{1\leq j\leq J_\ell} \delta_{\ell,j}^\top\,\mathsf{y}^{(\ell)} - h_\ell^*(\delta_{\ell,j}). \tag{4.8}$$

In particular, let $\{\mathsf{a}_{\ell,1},\dots,\mathsf{a}_{\ell,J_\ell}\} \subset \mathbb{R}^{M_\ell}$ be distinct elements in the relative interior of $\operatorname{dom} h_\ell$ and, for every $j \in \{1,\dots,J_\ell\}$, let $\delta_{\ell,j} \in \partial h_\ell(\mathsf{a}_{\ell,j})$. By the definition of subgradient, for every $j \in \{1,\dots,J_\ell\}$, the following inequality holds

$$(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad \delta_{\ell,j}^\top(\mathsf{y}^{(\ell)} - \mathsf{a}_{\ell,j}) + h_\ell(\mathsf{a}_{\ell,j}) \leq h_\ell(\mathsf{y}^{(\ell)}), \tag{4.9}$$

which implies that $h_\ell^*(\delta_{\ell,j}) = \delta_{\ell,j}^\top \mathsf{a}_{\ell,j} - h_\ell(\mathsf{a}_{\ell,j})$. Altogether, the latter and (4.8) show that a lower approximation of h_ℓ is given by

$$\widehat{h}_\ell(\mathsf{y}^{(\ell)}) = \max_{1\leq j\leq J_\ell} \delta_{\ell,j}^\top\,\mathsf{y}^{(\ell)} + \underbrace{h_\ell(\mathsf{a}_{\ell,j}) - \delta_{\ell,j}^\top \mathsf{a}_{\ell,j}}_{\mu_{\ell,j}}. \tag{4.10}$$

The above piecewise-affine function is completely defined by the *bundle* set

$$\mathcal{B}_\ell = \big\{(\mathsf{a}_{\ell,j}, \delta_{\ell,j}) \;\big|\; j \in \{1,\dots,J_\ell\}\big\}, \tag{4.11}$$

from which depends the tightness of the lower approximation in (4.4).

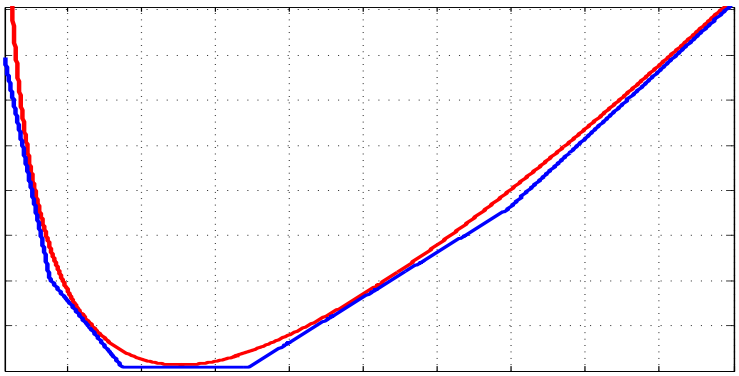

Figure 4.1 *A function (red line) and its piecewise-affine approximation (blue line).*

4.2.2 Asymptotic behavior

The validity of the proposed approximation can be analyzed by considering a sequence $(\widehat{h}^t)_{t\in\mathbb{N}}$ of progressively tighter approximations of the function h used in Problem (4.3). For every $t \in \mathbb{N}$, the function $\widehat{h}^t \in \Gamma_0(\mathbb{R}^M)$ is defined as

$$\widehat{h}^t(y) = \sum_{\ell=1}^{L} \widehat{h}^t_\ell(\mathsf{y}^{(\ell)}), \tag{4.12}$$

where

$$\widehat{h}^t_\ell(\mathsf{y}^{(\ell)}) = \max_{1\le j\le J^t_\ell} h_\ell(\mathsf{a}^t_{\ell,j}) + (\delta^t_{\ell,j})^\top(\mathsf{y}^{(\ell)} - \mathsf{a}^t_{\ell,j}). \tag{4.13}$$

The above piecewise-affine function is defined from the bundle set

$$\mathcal{B}^t_\ell = \big\{(\mathsf{a}^t_{\ell,j}, \delta^t_{\ell,j}) \;\big|\; j \in \{1,\dots,J^t_\ell\}\big\}, \tag{4.14}$$

where $\mathsf{a}^t_{\ell,j} \in \mathbb{R}^{M_\ell}$ and $\delta^t_{\ell,j} \in \partial h_\ell(\mathsf{a}^t_{\ell,j})$.

For every $t \in \mathbb{N}$, $\widehat{C}^t = \operatorname{lev}_{\le\eta} \widehat{h}^t$ is an outer approximation of $C = \operatorname{lev}_{\le\eta} h$. Under appropriate assumptions, it is possible to prove that such an approximation asymptotically leads to a solution to Problem (4.3), as explained below.

Assumption 4.2.1.

(i). *There exists* $\overline{x} \in \mathbb{R}^N$ *such that* $T_1\overline{x} \in \operatorname{dom} f_1, \dots, T_R\overline{x} \in \operatorname{dom} f_R, F\overline{x} \in C$.

(ii). $\sum_{r=1}^R f_r \circ T_r$ *is coercive or* $F^{-1}(\widehat{C}^0)$ *is bounded.*

(iii). $\mathcal{B}^t_\ell \subset \mathcal{B}^{t+1}_\ell$ *for every* $t \in \mathbb{N}$ *and for every* $\ell \in \{1,\dots,L\}$.

(iv). *For every* $\ell \in \{1,\dots,L\}$ *and for every* $x \in \mathbb{R}^N$ *such that* $F_\ell x \in \operatorname{dom} h_\ell$,

$$\lim_{t\to+\infty} \widehat{h}^t_\ell(F_\ell x) = h_\ell(F_\ell x). \tag{4.15}$$

(v). $\bigcap_{r=1}^R T_r^{-1}(\operatorname{dom} f_r) \subset F^{-1}(\operatorname{dom} h)$.

Proposition 4.2.2. *Under Assumption 4.2.1, the following points hold.*

(i). For every $t \in \mathbb{N}$, there exists a minimizer $\widehat{x}^t$ to $\sum_{r=1}^R f_r \circ T_r$ over $F^{-1}(\widehat{C}^t)$.

(ii). Let $(\widehat{x}^t)_{t\in\mathbb{N}}$ be a sequence of minimizers. Then, $(\widehat{x}^t)_{t\in\mathbb{N}}$ is a bounded sequence, any of its cluster points is a solution $\widehat{x}$ to Problem (4.3), and $\big(\sum_{r=1}^R f_r(T_r\widehat{x}^t)\big)_{t\in\mathbb{N}}$ is an increasing sequence converging to $\sum_{r=1}^R f_r(T_r\widehat{x})$.

(iii). If Problem (4.3) has a unique solution $\widehat{x}$, then $(\widehat{x}^t)_{t\in\mathbb{N}}$ converges to $\widehat{x}$.

Proof. Problem 4.3 is equivalent to minimize $\varphi = \sum_{r=1}^R f_r \circ T_r + \iota_C \circ F$. For every $t \in \mathbb{N}$, let $\widehat{\varphi}^t = \sum_{r=1}^R f_r \circ T_r + \iota_{\widehat{C}^t} \circ F$. Due to Assumption 4.2.1(iii),

$$\widehat{h}^t \leq \widehat{h}^{t+1} \leq h, \tag{4.16}$$

which implies that $C \subset \widehat{C}^{t+1} \subset \widehat{C}^t$, and thus

$$\widehat{\varphi}^t \leq \widehat{\varphi}^{t+1} \leq \varphi. \tag{4.17}$$

From this inequality and Assumption 4.2.1(i), it can be deduced that the functions φ and $(\widehat{\varphi}^t)_{t\in\mathbb{N}}$ are proper, thus they belong to $\Gamma_0(\mathbb{R}^N)$.

In addition, due to Assumption 4.2.1(ii), φ^0 is coercive (i.e., level-bounded), and thus the sequence $(\widehat{\varphi}^t)_{t\in\mathbb{N}}$ is eventually level-bounded (see [192, Ex. 7.32]). Besides, $(\widehat{\varphi}^t)_{t\in\mathbb{N}}$ converges pointwise to φ according to Assumptions 4.2.1(iv)-(v). As $(\widehat{\varphi}^t)_{t\in\mathbb{N}}$ is an increasing sequence, Proposition 7.4(d) in [192] shows that $(\widehat{\varphi}^t)_{t\in\mathbb{N}}$ epi-converges to φ. It can then be deduced from [192, Theorem 7.33] that, for every $t \in \mathbb{N}$, $\operatorname{Argmin} \widehat{\varphi}^t \neq \emptyset$, which proves (i). Theorem 7.33 in [192] also allows us to claim that any sequence of minimizers $(\widehat{x}^t)_{t\in\mathbb{N}}$ is bounded, its cluster points belong to $\operatorname{Argmin} \varphi$, and $\widehat{\varphi}^t(\widehat{x}^t) \to \inf \varphi$ as $t \to +\infty$. The fact that $(\widehat{\varphi}^t(\widehat{x}^t))_{t\in\mathbb{N}}$ is an increasing sequence follows from (4.17), which completes the proof of (ii). The point (iii) is a direct consequence of (ii). □

A natural question in this context is how the sets $(\mathcal{B}_\ell^t)_{t\in\mathbb{N}}$ can be chosen so that Assumption 4.2.1(iv) is satisfied. The following result provides an answer.

If $\operatorname{dom} h_\ell^*$ is open, the continuity property is clearly satisfied.

Proposition 4.2.3. *Let $\ell \in \{1, \dots, L\}$. Assume that h_ℓ^* is continuous relative to its domain and that, for every $\delta_\ell \in \operatorname{dom} h_\ell^*$, there exists a sequence $(\mathsf{a}_\ell^t, \delta_\ell^t)_{t\in\mathbb{N}}$ such that $(\mathsf{a}_\ell^t, \delta_\ell^t) \in \mathcal{B}_\ell^t$ for each $t \in \mathbb{N}$, and*

$$\lim_{t\to+\infty} \delta_\ell^t = \delta_\ell. \tag{4.18}$$

Then,

$$(\forall \mathsf{y}^{(\ell)} \in \operatorname{dom} h_\ell) \qquad \lim_{t\to+\infty} \widehat{h}_\ell^t(\mathsf{y}^{(\ell)}) = h_\ell(\mathsf{y}^{(\ell)}). \tag{4.19}$$

Proof. We proceed similarly to the proofs in [74, 158]. Let $\mathsf{y}^{(\ell)} \in \operatorname{dom} h_\ell$. Then,

$$h_\ell(\mathsf{y}^{(\ell)}) = \sup_{\delta_\ell \in \operatorname{dom} h_\ell^*} (\delta^{(\ell)})^\top \mathsf{y}^{(\ell)} - h_\ell^*(\delta_\ell) < +\infty. \tag{4.20}$$

For any given positive real ε, there thus exists $\chi_\varepsilon^{(\ell)} \in \operatorname{dom} h_\ell^*$ such that

$$h_\ell(\mathsf{y}^{(\ell)}) \leq (\chi_\varepsilon^{(\ell)})^\top \mathsf{y}^{(\ell)} - h_\ell^*(\chi_\varepsilon^{(\ell)}) + \frac{\varepsilon}{3}. \tag{4.21}$$

Since h_ℓ^* is continuous relative to its domain, there exists a neighbourhood $\mathcal{N}(\chi_\varepsilon^{(\ell)})$ of $\chi_\varepsilon^{(\ell)}$ such that

$$\big(\forall \delta_\ell \in \mathcal{N}(\chi_\varepsilon^{(\ell)}) \cap \operatorname{dom} h_\ell^*\big) \qquad h_\ell^*(\delta_\ell) \leq h_\ell^*(\chi_\varepsilon^{(\ell)}) + \frac{\varepsilon}{3}. \tag{4.22}$$

In addition, we know that there exists a sequence $(\mathsf{a}_\ell^t, \delta_\ell^t)_{t\in\mathbb{N}}$ such that, for every $t \in \mathbb{N}$, $(\mathsf{a}_\ell^t, \delta_\ell^t) \in \mathcal{B}_\ell^t$ and $\lim_{t\to+\infty} \delta_\ell^t = \chi_\varepsilon^{(\ell)}$. Hence, there exists $t_\varepsilon \in \mathbb{N}$ such that, for every $t \geq t_\varepsilon$, $\delta_\ell^t \in \mathcal{N}(\chi_\varepsilon^{(\ell)}) \cap \operatorname{dom} h_\ell^*$ and

$$(\chi_\varepsilon^{(\ell)} - \delta_\ell^t)^\top \mathsf{y}^{(\ell)} \leq \frac{\varepsilon}{3}. \tag{4.23}$$

Now, by noticing that

$$\widehat{h}_\ell^t(\mathsf{y}^{(\ell)}) \geq (\delta_\ell^t)^\top \mathsf{y}^{(\ell)} - h_\ell^*(\delta_\ell^t), \tag{4.24}$$

(4.21), (4.22) and (4.23) yield, for every $t \geq t_\varepsilon$,

$$\begin{aligned} 0 &\leq h_\ell(\mathsf{y}^{(\ell)}) - \widehat{h}_\ell^t(\mathsf{y}^{(\ell)}) \\ &\leq (\chi_\varepsilon^{(\ell)} - \delta_\ell^t)^\top \mathsf{y}^{(\ell)} + h_\ell^*(\delta_\ell^t) - h_\ell^*(\chi_\varepsilon^{(\ell)}) + \frac{\varepsilon}{3} \leq \varepsilon. \end{aligned} \tag{4.25}$$

This shows that $\big(\widehat{h}_\ell^t(\mathsf{y}^{(\ell)})\big)_{t\in\mathbb{N}}$ converges to $h_\ell(\mathsf{y}^{(\ell)})$. □

To illustrate these results, the Kullback-Leibler divergence is now considered.

Example 4.2.4. *For every $\ell \in \{1, \ldots, L\}$, let $M_\ell = 1$, $z^{(\ell)} \in [0, +\infty[$, and*

$$(\forall y^{(\ell)} \in \mathbb{R}) \qquad h_\ell^{\mathrm{KL}}(y^{(\ell)}) = \begin{cases} -z^{(\ell)} \ln(y^{(\ell)}) + \alpha\, y^{(\ell)}, & \text{if } y^{(\ell)} > 0 \text{ and } z^{(\ell)} > 0, \\ \alpha\, y^{(\ell)}, & \text{if } y^{(\ell)} \geq 0 \text{ and } z^{(\ell)} = 0, \\ +\infty, & \text{otherwise,} \end{cases} \tag{4.26}$$

where $\alpha > 0$. Then,

$$(\forall \delta_\ell \in \mathbb{R}) \qquad (h_\ell^{\mathrm{KL}})^*(\delta_\ell) = \begin{cases} -z^{(\ell)} \ln\left(\dfrac{\alpha - \delta_\ell}{z_\ell}\right) - z^{(\ell)}, & \text{if } \delta_\ell < \alpha \text{ and } z^{(\ell)} > 0, \\ 0 & \text{if } \delta_\ell = \alpha \text{ and } z^{(\ell)} = 0, \\ +\infty, & \text{otherwise.} \end{cases} \tag{4.27}$$

For every $t \in \mathbb{N}$, set $J_\ell^t = (t+1)2^t$ and choose

$$\mathcal{B}_\ell^t = \left\{ \left(\frac{z^{(\ell)}\, 2^t}{j}, \alpha_\ell - \frac{j}{2^t} \right) \;\middle|\; j \in \{1, \ldots, J_\ell^t\} \right\}. \tag{4.28}$$

Then, Assumption 4.2.1(iii) holds and Proposition 4.2.3 shows that Assumption 4.2.1(iv) is also satisfied.

Note that these asymptotic properties are more of theoretical interest than of practical one. Indeed, the applications considered in Section 4.4 will show that satisfactory results are obtained for small values of the bundle size.

4.3 Optimization method

The key point in the proposed method is that one solves the approximated problem (4.6) in order to find a solution to the original problem (4.3). This approach is motivated by the fact that the approximated problem is better suited for the epigraphical splitting. Indeed, the piecewise-affine function $\widehat{h}$ in (4.10) can be expressed as the max function composed with a linear operator, namely

$$(\forall \mathsf{y}^{(\ell)} \in \mathbb{R}^{M_\ell}) \qquad \widehat{h}_\ell(\mathsf{y}^{(\ell)}) = h_\ell^{\max}(\Delta_\ell\, \mathsf{y}^{(\ell)}), \tag{4.29}$$

where the matrix Δ_ℓ is built from the bundled subgradients $(\delta_{\ell,j})_{1\leq j\leq J_\ell}$ of h_ℓ

$$\Delta_\ell = \begin{bmatrix} \delta_{\ell,1}^\top \\ \vdots \\ \delta_{\ell,J_\ell}^\top \end{bmatrix} \in \mathbb{R}^{J_\ell \times M_\ell}, \tag{4.30}$$

and the function $h_\ell^{\max}$ is defined as the component-wise maximum

$$(\forall \mathsf{u}^{(\ell)} = (u^{(\ell,j)})_{1\leq j\leq J_\ell} \in \mathbb{R}^{J_\ell}) \qquad h_\ell^{\max}(\mathsf{u}^{(\ell)}) = \max_{1\leq j\leq J_\ell}\; u^{(\ell,j)} + \mu_{\ell,j}. \tag{4.31}$$

As a result, the epigraphical decomposition of Problem (4.6) yields

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^L}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\, x) \quad \text{s.t.} \quad \begin{cases} (\forall \ell \in \{1,\dots,L\}) \quad h_\ell^{\max}\,(\Delta_\ell F_\ell\, x) \leq \zeta^{(\ell)}, \\ \displaystyle\sum_{\ell=1}^{L} \zeta^{(\ell)} \leq \eta. \end{cases} \tag{4.32}$$

By setting $F = [F_1^\top \;\dots\; F_L^\top]^\top \in \mathbb{R}^{M\times N}$ and $\Delta = \operatorname{diag}(\Delta_1,\dots,\Delta_L) \in \mathbb{R}^{J\times M}$ (with $M = M_1 + \cdots + M_L$ and $J = J_1 + \cdots + J_L$), the above problem can be more conveniently rewritten as follows

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^L}{\text{minimize}} \quad \sum_{r=1}^{R} f_r(T_r\, x) \quad \text{s.t.} \quad \begin{cases} (\Delta F x, \zeta) \in E, \\ \zeta \in V, \end{cases} \tag{4.33}$$

where

$$E = \big\{(u,\zeta) \in \mathbb{R}^J \times \mathbb{R}^L \;\big|\; (\forall \ell \in \{1,\dots,L\}) \quad (\mathsf{u}^{(\ell)}, \zeta^{(\ell)}) \in \operatorname{epi} h_\ell^{\max}\big\}, \tag{4.34}$$

$$V = \big\{\zeta \in \mathbb{R}^L \;\big|\; 1_L^\top \zeta \leq \eta\big\}, \tag{4.35}$$

with the generic vector $u \in \mathbb{R}^J$ being block-decomposed as follows

$$u = \Big[\underbrace{\big(\mathsf{u}^{(1)}\big)^\top}_{\text{size } J_1}, \dots, \underbrace{\big(\mathsf{u}^{(L)}\big)^\top}_{\text{size } J_L}\Big]^\top \in \mathbb{R}^{J=J_1+\cdots+J_L}. \tag{4.36}$$

The undoubtedly advantage of the above reformulation is that the projection onto E can be computed in closed form (see Section 2.3.3), opening the way to efficient implementations of Problem (4.33) based on primal-dual proximal methods. Some practical examples will be presented in the next section.

4.4 Numerical results

The present section illustrates the validity of the proposed approximation method by means of two practical applications, such as image recovery in the presence of Poisson noise, and logistic regression. It is worth emphasizing that the primarily aim of the numerical analysis presented hereafter is to compare the proposed approach with "exact" state-of-the-art solutions [137, 207], so as to assess the impact of the approximation over the quality of obtained results.

4.4.1 Image recovery with Poisson noise

Digital imaging sensors are inherently affected by *Poisson noise*, a signal-dependent form of uncertainty originated from the quantized nature of light. This phenomenon is very intense in applications such as emission tomography, microscopy, and astronomy. As a result, it is common to model the degradation process with a linear operator followed by a Poisson noise, namely

$$z = \mathcal{P}_\alpha(A\overline{x}), \tag{4.37}$$

where $\overline{x} \in \mathbb{R}^N$, $A \in \mathbb{R}^{K\times N}$, and $z \in [0, +\infty[^K$ are defined as in (3.1), while $\mathcal{P}_\alpha$ denotes a realization of Poisson noise with scaling parameter $\alpha > 0$.

A possible approach to recover $\overline{x}$ from z is to formulate a convex optimization problem in which the TV regularization in (3.4) is optimized under a constraint involving the KL divergence in (4.26), yielding [207]

$$\underset{x\in\mathbb{R}^N}{\text{minimize}} \quad \mathrm{TV}_2(x) \quad \text{s. t.} \quad \sum_{\ell=1}^{K} h_\ell^{\mathrm{KL}}\Big((Ax)^{(\ell)}\Big) \leq K/2, \tag{4.38}$$

where the bound $K/2$ follows from the statistical properties of Poisson noise [238]. According to the optimization method detailed in Section 4.3, the above problem can be approximated through an instance of Problem (4.33), namely

$$\underset{(x,\zeta)\in\mathbb{R}^N\times\mathbb{R}^K}{\text{minimize}} \quad \|Fx\|_{1,2} \quad \text{s. t.} \quad \begin{cases} (\Delta Ax, \zeta) \in E, \\ \zeta \in V, \end{cases} \tag{4.39}$$

where $\|\cdot\|_{1,2}$ is the $\ell_{1,2}$-norm defined in (3.3), and $F \in \mathbb{R}^{2N\times N}$ is the gradient operator defined in (3.5). The iterations of PDFB [69, 215] associated with the approximated problem (4.39) are given in Algorithm 4.1. For sake of comparison, the iterations of PDFB associated to the exact problem (4.38) are given in Algorithm 4.2, where the projection onto $D = \big\{y \in \mathbb{R}^K \;\big|\; \sum_{\ell=1}^{K} h_\ell^{\mathrm{KL}}(y^{(\ell)}) \leq K/2\big\}$ is computed with a specific iterative method [207], as detailed in Section 1.3.3.

Figure 4.2 illustrates an example of image degraded by the model in (4.37) using a uniform blur of size 3×3 and $\alpha = 1$, along with the images recovered by solving the exact problem in (4.38), and the approximated one in (4.39) for several sizes of the bundle set indicated in (4.28). The SNR indexes show that satisfactory results can be obtained with small bundle sets, while the execution times indicate the efficiency of the proposed optimization method w.r.t. the approach based on the direct computation of projections via inner iterations [207]. In this experiment, the stopping criterion was set to $\|x^{[i+1]} - x^{[i]}\| \leq 10^{-4}\|x^{[i]}\|$, the algorithms were developed in MATLAB R2011b,

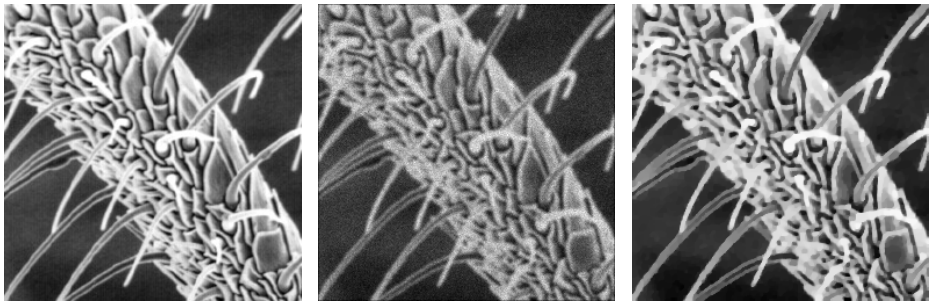

(a) Original image. (b) Degraded. (c) Exact [207]. SNR–SSIM: 21.98–0.869. Time: 16.35 s.

(d) Approx with $J_k = 9$. SNR–SSIM: 21.80–0.861. Time: 44.05 s. (e) Approx with $J_\ell = 13$. SNR–SSIM: **22.04–0.862**. Time: 59.47 s. (f) Approx with $J_\ell = 17$. SNR–SSIM: 22.04–0.862. Time: 70.25 s.

Figure 4.2 *SNR (dB) – SSIM indexes for the recovery of a grayscale image degraded with a uniform blur of size* 3×3 *and a Poisson noise with scaling* $\alpha = 1$.

Algorithm 4.1 *FBPD for Problem* (4.39)

INITIALIZATION

choose $\big(x^{[0]}, y^{[0]}, w^{[0]}\big) \in \mathbb{R}^N \times \mathbb{R}^J \times \mathbb{R}^{2N}$

choose $\big(\zeta^{[0]}, \xi^{[0]}\big) \in \mathbb{R}^K \times \mathbb{R}^K$

set $\tau > 0$ and $\sigma > 0$ such that

$$\tau\sigma\big(\|F\|^2 + \max\{\|\Delta A\|^2, 1\}\big) \leq 1$$

FOR $i = 0, 1, \ldots$

$$x^{[i+1]} = x^{[i]} - \tau\big(A^\top \Delta^\top y^{[i]} + F^\top w^{[i]}\big)$$
$$\zeta^{[i+1]} = P_V\left(\zeta^{[i]} - \gamma\,\xi^{[i]}\right)$$
$$\widehat{y}^{[i]} = y^{[i]} + \sigma A\big(2x^{[i+1]} - x^{[i]}\big)$$
$$\widehat{\xi}^{[i]} = \xi^{[i]} + \sigma\big(2\zeta^{[i+1]} - \zeta^{[i]}\big)$$
$$\widehat{w}^{[i]} = w^{[i]} + \sigma F\big(2x^{[i+1]} - x^{[i]}\big)$$
$$(y^{[i+1]}, \xi^{[i+1]}) = (\widehat{y}^{[i]}, \widehat{\xi}^{[i]}) - \sigma P_E\big(\tfrac{\widehat{y}^{[i]}}{\sigma}, \tfrac{\widehat{\xi}^{[i]}}{\sigma}\big)$$
$$w^{[i+1]} = \widehat{w}^{[i]} - \sigma \operatorname{prox}_{\frac{1}{\sigma}\|\cdot\|_{1,2}}\big(\widehat{w}^{[i]}/\sigma\big)$$

Algorithm 4.2 *FBPD for Problem* (4.38)

INITIALIZATION

choose $x^{[0]} \in \mathbb{R}^N$

choose $\big(y^{[0]}, w^{[0]}\big) \in \mathbb{R}^J \times \mathbb{R}^{2N}$

set $\tau > 0$ and $\sigma > 0$ such that

$$\tau\sigma\big(\|F\|^2 + \|A\|^2\big) \leq 1$$

FOR $i = 0, 1, \ldots$

$$x^{[i+1]} = x^{[i]} - \tau\big(A^\top y^{[i]} + F^\top w^{[i]}\big)$$
$$\widehat{y}^{[i]} = y^{[i]} + \sigma A\big(2x^{[i+1]} - x^{[i]}\big)$$
$$\widehat{w}^{[i]} = w^{[i]} + \sigma F\big(2x^{[i+1]} - x^{[i]}\big)$$
$$y^{[i+1]} = \widehat{y}^{[i]} - \sigma P_D\big(\widehat{y}^{[i]}/\sigma\big)$$
$$w^{[i+1]} = \widehat{w}^{[i]} - \sigma \operatorname{prox}_{\frac{1}{\sigma}\|\cdot\|_{1,2}}\big(\widehat{w}^{[i]}/\sigma\big)$$

4.4.2 Classification with logistic regression

Classification is the task of identifying the class $k \in \{1, \dots, K\}$ to which belongs an observation $u \in \mathbb{R}^N$ (e.g., a signal, image, or graph). To do so, one can select the class that maximizes the correlation with some "references", namely

$$d(u) = \operatorname*{argmax}_{k \in \{1,\dots,K\}} \phi(u)^\top x^{(k)}, \tag{4.40}$$

where $\phi \colon \mathbb{R}^N \mapsto \mathbb{R}^M$ denotes the mapping into an arbitrary feature space, and $x = [(x^{(1)})^\top \ \dots \ (x^{(K)})^\top]^\top \in \mathbb{R}^{MK}$ denotes the vector of references that (in supervised learning) have to be estimated from a given set of input-output pairs

$$S = \big\{(u_\ell, z_\ell) \in \mathbb{R}^N \times \{1, \dots, K\} \;\big|\; \ell \in \{1, \dots, L\}\big\}. \tag{4.41}$$

Finding the vector $\overline{x}$ that allows (4.40) to correctly predict *all* the training samples is an ill-posed problem, as there is no guarantee that such a vector exists. Logistic regression overcomes this issue by resorting to a *maximum a posteriori probability* (MAP) estimation of the vector $\overline{x}$. To this end, the likelihood of the training samples is modeled by the normalized exponential distribution

$$(\forall \ell \in \{1, \dots, L\}) \qquad \Pr\big(d(u_\ell) = z_\ell \mid x\big) = \frac{\exp\big(\phi(u_\ell)^\top x^{(z_\ell)}\big)}{\sum_{k=1}^K \exp\big(\phi(u_\ell)^\top x^{(k)}\big)}. \tag{4.42}$$

The MAP approach thus amounts to minimizing the anti-log of the above likelihood and a Gaussian prior distribution (scaled by a factor $\lambda > 0$), yielding

$$\operatorname*{minimize}_{x \in \mathbb{R}^{MK}} \quad \lambda \|x\|_2^2 + \sum_{\ell=1}^L \log\Big(1 + \sum_{k \neq z_\ell} e^{\phi(u_\ell)^\top (x^{(k)} - x^{(z_\ell)})}\Big). \tag{4.43}$$

Despite the popularity of quadratic regularization, sparsity-inducing penalties have recently attracted much attention in machine learning, as they perform an implicit "feature selection" by shrinking small coefficients to zero [9, 23, 137, 189, 214, 218]. A possible formulation of the sparse logistic regression is

$$\operatorname*{minimize}_{x \in \mathbb{R}^{MK}} \quad \sum_{k=1}^K \|x^{(k)}\|_{1,p} \quad \text{s. t.} \quad \sum_{\ell=1}^L h_\ell^{\log}(F_\ell\, x) \leq \eta, \tag{4.44}$$

where $\eta \geq 0$, $F_\ell\, x = [\phi(u_\ell)^\top (x^{(k)} - x^{(z_\ell)})]_{1 \leq k \leq K} \in \mathbb{R}^{K \times MK}$, and

$$(\forall \mathsf{y}^{(\ell)} = (y^{(\ell,k)})_{1 \leq k \leq K} \in \mathbb{R}^K) \qquad h_\ell^{\log}(\mathsf{y}^{(\ell)}) = \log\Big(\sum_{k=1}^K \exp\big(y^{(\ell,k)}\big)\Big). \tag{4.45}$$

According to the optimization method detailed in Section 4.3, the above problem can be approximated through an instance of Problem (4.33), namely

$$\operatorname*{minimize}_{(x,\zeta) \in \mathbb{R}^{MK} \times \mathbb{R}^L} \quad f(x) \quad \text{s. t.} \quad \begin{cases} (\Delta F x, \zeta) \in E, \\ \zeta \in V, \end{cases} \tag{4.46}$$

where $f \in \Gamma_0(\mathbb{R}^{MK})$ denotes the regularization term, and $F = [F_1^\top \dots F_L^\top]^\top$. The iterations of PDFB associated to Problem (4.46) are given in Algorithm 4.3.

For sake of comparison, the regularized version of Problem (4.44) is considered as well, which amounts to

$$\underset{x\in\mathbb{R}^{MK}}{\text{minimize}} \quad f(x) + \lambda \sum_{\ell=1}^{L} h_\ell^{\log}(F_\ell\, x), \tag{4.47}$$

with $\lambda > 0$. The above problem can be solved exactly through the iterations of forward-backward splitting summarized in Algorithm 4.4, as the function $h_\ell^{\log}$ is differentiable with 1-Lipschitz continuous gradient [24].

Figure 4.3 illustrates a database containing a large number of 28×28 images ($N = 784$) displaying handwritten digits ($K = 10$) organized in 60000 training images and 10000 test images [142]. This database is used to evaluate the performance of logistic regression w.r.t. the exact regularized problem (4.47), and the approximated problem (4.46) for several sizes of the bundle set (built by uniformly sampling the hypercube $[-3, 3]^K$). In both cases, the mapping ϕ is given by Kernel PCA [197] with a Gaussian kernel (thus $M = N$).

Table 4.1 collects the classification errors obtained by running the classifiers on the 10000 test images, after a learning performed on training sets of different size $L \in \{5K, 10K, 20K, 50K\}$ and with different values of the parameter η. For the regularization, we used the $\boldsymbol{\ell}_{1,\infty}$-norm computed by dividing each vector $(x^{(k)})_{1\leq k\leq K}$ in blocks of size 10. The results indicate that the proposed approximation is slightly sensitive to the size of the bundle set, and that even a coarse approximation allows one to achieve a performance close to the ideal case, within an error of 0.5%. It is worth emphasizing that the proposed approach aims at directly solving Problem (4.44), while state-of-the-art methods only focus on the regularized problem (4.47), which presents no challenge from an optimization standpoint.

4.5 Conclusion

In this chapter, we have proposed a technique for outer approximating the sublevel set of a block-separable function. We have then derived an optimization method for dealing with such an approximated constraint through the epigraphical splitting. To illustrate the validity of our approach, we have tackled constraints involving the Kullback-Leibler divergence and the logistic loss. In the first case, the results obtained with the proposed approximation are equivalent to those achieved with the exact approach. In the second case, the proposed approximation allows us to derive an algorithm for solving the sparse logistic regression in constrained form, and the comparison with the regularized version shows that our approach achieves a performance close to the exact case.

In the next part, we will employ the epigraphical splitting to design competitive algorithms in the context of multicomponent image recovery, classification with sparse support vector machines, and PRNU-based image forgery detection.

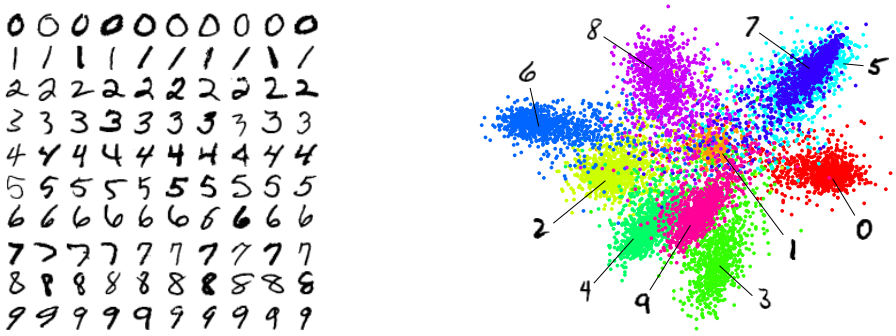

(a) Training images. (b) Scatter plot in a feature space.

Figure 4.3 *MNIST database.*

Algorithm 4.3 *FBPD algorithm [69, 215] specialized to Problem* (4.46)

INITIALIZATION

choose $\big(x^{[0]}, \zeta^{[0]}\big) \in \mathbb{R}^N \times \mathbb{R}^L$

choose $\big(y^{[0]}, \xi^{[0]}\big) \in \mathbb{R}^J \times \mathbb{R}^L$

set $\tau > 0$ and $\sigma > 0$ such that $\tau\sigma \max\{\|\Delta F\|^2, 1\} \leq 1$

FOR $i = 0, 1, \ldots$

$$x^{[i+1]} = \operatorname{prox}_{\tau f}\left(x^{[i]} - \tau F^\top \Delta^\top y^{[i]}\right)$$
$$\zeta^{[i+1]} = P_V\left(\zeta^{[i]} - \gamma\, \xi^{[i]}\right)$$
$$\widehat{y}^{[i]} = y^{[i]} + \sigma F\big(2x^{[i+1]} - x^{[i]}\big)$$
$$\widehat{\xi}^{[i]} = \xi^{[i]} + \sigma\big(2\zeta^{[i+1]} - \zeta^{[i]}\big)$$
$$(y^{[i+1]}, \xi^{[i+1]}) = (\widehat{y}^{[i]}, \widehat{\xi}^{[i]}) - \sigma P_E\big(\widehat{y}^{[i]}/\sigma, \widehat{\xi}^{[i]}/\sigma\big)$$

Algorithm 4.4 *Forward-backward splitting [68] specialized to Problem* (4.47)

INITIALIZATION

choose $x^{[0]} \in \mathbb{R}^N$

set $\gamma \in \,]0, 2[$

FOR $i = 0, 1, \ldots$

$$x^{[i+1]} = \operatorname{prox}_{\gamma f}\left(x^{[i]} - \gamma \sum_{\ell=1}^{L} F_\ell^\top \nabla h_\ell^{\log}(F_\ell\, x^{[i]})\right)$$

Table 4.1 *Classification errors (in the case $K = 3$) obtained with the lower-approximated logistic regression for several values of L, η and $J_\ell = (2m+1)^{K-1}$.*

	η	$g = \|\|\cdot\|\|_{1,\infty}$				
		$m=3$	$m=5$	$m=7$	$m=9$	exact
$L/K = 5$	$0.01L$	3.37 %	3.40 %	3.37 %	3.43 %	2.76 %
	$0.05L$	3.08 %	3.53 %	3.24 %	3.21 %	
	$0.1L$	3.11 %	2.83 %	3.08 %	2.92 %	
	$0.5L$	2.45 %	3.08 %	3.11 %	**2.57** %	
	$0.7L$	3.15 %	2.64 %	2.80 %	2.80 %	
	$1.0L$	6.04 %	6.32 %	6.70 %	6.93 %	
$L/K = 10$	$0.01L$	3.05 %	2.96 %	2.89 %	2.92 %	2.51 %
	$0.05L$	2.48 %	2.48 %	2.86 %	2.48 %	
	$0.1L$	2.67 %	2.64 %	2.80 %	2.83 %	
	$0.5L$	2.32 %	2.67 %	2.51 %	**2.29** %	
	$0.7L$	2.80 %	2.70 %	2.41 %	2.41 %	
	$1.0L$	6.77 %	4.26 %	3.62 %	3.56 %	
$L/K = 20$	$0.005L$	1.78 %	2.22 %	2.03 %	2.07 %	1.97 %
	$0.01L$	2.35 %	2.07 %	2.00 %	2.07 %	
	$0.05L$	1.68 %	2.03 %	1.97 %	2.05 %	
	$0.1L$	2.07 %	2.07 %	1.94 %	**2.03** %	
	$0.5L$	2.32 %	2.76 %	2.67 %	2.70 %	
	$0.7L$	3.53 %	2.99 %	3.02 %	2.92 %	
	$1.0L$	8.64 %	6.23 %	7.05 %	6.55 %	
$L/K = 50$	$0.005L$	1.56 %	1.08 %	1.08 %	**1.14** %	1.49 %
	$0.01L$	1.56 %	1.43 %	1.18 %	1.53 %	
	$0.05L$	1.84 %	1.43 %	1.53 %	1.56 %	
	$0.1L$	2.03 %	1.84 %	1.81 %	1.84 %	
	$0.5L$	2.89 %	2.86 %	2.76 %	2.80 %	
	$0.7L$	4.26 %	3.37 %	3.37 %	3.34 %	
	$1.0L$	12.81 %	10.84 %	12.20 %	11.63 %	

Part II

Applications

Knowledge without application is like a book that is never read.

Christopher Crawford

Chapter 5
Nonlocal structure tensor

The main focus of this chapter is multicomponent image recovery from degraded observations, for which it is of paramount importance to exploit the intrinsic correlations along spatial and spectral dimensions. To this end, we propose a variational approach based on the introduction of a "nonlocal structure tensor" regularity constraint, and we show how to solve it practically with the proposed epigraphical splitting. Experiments on color, multispectral and hyperspectral images demonstrate the validity of our approach.

5.1 Multicomponent image recovery problems

Multicomponent images consist of several spatial maps acquired simultaneously from a scene. Well-known examples are color images, which are composed of red, green, and blue components, or spectral images, which divide the electromagnetic spectrum into many components that represent the light intensity across a number of wavelengths. Multicomponent images are often degraded by blur and noise arising from sensor imprecisions or physical limitations, such as aperture effects, motion, or atmospheric phenomena. Additionally, a decimation modelled by a sparse or random matrix can be encountered in several applications, such as compressive sensing [107, 190, 193], inpainting [96, 157, 199], or super-resolution [240]. As a consequence, the standard imaging model consists of a blurring operator [176] followed by a decimation and some kind of noise.

The quality of results obtained through a variational approach strongly depends on the ability to model the regularity present in images. Since natural images are often piecewise smooth, popular regularization models, such as total variation (TV) and its generalizations, tend to penalize the image gradient (see Section 3.1). The extension of TV-based models to multicomponent images is however nontrivial. A first approach consists of computing TV channel-by-channel and then summing up the resulting smoothness measures [6, 18, 89, 236]. Since there is no coupling of the components, this approach may potentially lead to component smearing and loss of edges across components. An alternative way is to process the components jointly, so as to better reveal details and features that are not visible in each of the components considered separately. This approach naturally arises when the gradient of a multicomponent image is thought of as a *structure tensor* (ST) [28, 84, 108, 129, 145, 196, 201, 210, 222], namely a matrix that summarizes the prevailing direction of the gradient.

The principle of ST-based regularization is to penalize the eigenvalues of the ST matrix, in order to smooth in the direction of minimal change [18, 108]. The first contribution of this chapter is the introduction of a new regularization for multicomponent images that penalizes the eigenvalues of a nonlocal extension of the ST matrix, so as to combine the strengths of ST and nonlocal TV (ST-NLTV).

5.1.1 Degradation model

The R-component image of interest is denoted by $\overline{\mathrm{x}} = (\overline{x}_1, \ldots, \overline{x}_R) \in (\mathbb{R}^N)^R$, where each component corresponds to an image channel of size $N = N_1 \times N_2$. The degradation process of a multicomponent image can be modeled as

$$\mathrm{z} = \mathcal{B}(\mathrm{A}\overline{\mathrm{x}}), \tag{5.1}$$

where the linear operator $A\colon (\mathbb{R}^N)^R \mapsto (\mathbb{R}^K)^S$ is defined as

$$A = \begin{bmatrix} A_{1,1} & \ldots & A_{1,R} \\ \vdots & \ddots & \vdots \\ A_{S,1} & \ldots & A_{S,R} \end{bmatrix} \tag{5.2}$$

where $\mathrm{A}_{s,r} \in \mathbb{R}^{K\times N}$ for every $(s,r) \in \{1,\ldots,S\} \times \{1,\ldots,R\}$, the operator $\mathcal{B}\colon (\mathbb{R}^K)^S \to (\mathbb{R}^K)^S$ models the effect of a noise, and $\mathrm{z} = (z_1, \ldots, z_S) \in (\mathbb{R}^K)^S$ denotes the degraded observations. Interestingly, the above model can be specialized to the following applications.

(i). **Compressive sensing** [107]. In this scenario, z denotes the compressed multicomponent image, and $\overline{\mathrm{x}}$ is the multicomponent image to be reconstructed. The operator A is a block-diagonal matrix with $S = R$. For every $r \in \{1,\ldots,R\}$, the non-zero block $\mathrm{A}_{r,r}$ is a random measurement matrix $\mathrm{D}_r \in \mathbb{R}^{K\times N}$ with $K \ll N$. The noise is assumed to be a zero-mean white Gaussian additive noise. This leads to the following degradation model:

$$(\forall r \in \{1,\ldots,R\}) \qquad z_r = \mathrm{D}_r\, \overline{x}_r + \varepsilon_r \tag{5.3}$$

where $\varepsilon_r \sim \mathcal{N}(0, \sigma^2\, \mathrm{Id}_K)$.

(ii). **Super-resolution** [240]. In this scenario, z denotes B multicomponent images at low-resolution, and $\overline{\mathrm{x}}$ is the (high-resolution) multicomponent image to be recovered. The operator A is a block-diagonal matrix with $S = BR$. For every $r \in \{1,\ldots,R\}$ and $b \in \{1,\ldots,B\}$, the non-zero block $\mathrm{A}_{B(r-1)+b,r}$ is a composition of a warp matrix $\mathrm{W}_r \in \mathbb{R}^{N\times N}$, a blur operator $\mathrm{T} \in \mathbb{R}^{N\times N}$, and a downsampling matrix $\mathrm{D}_b \in \mathbb{R}^{K\times N}$ with $K < N$. The noise is assumed to be a zero-mean white Gaussian additive noise. As a result, there are B noisy decimated versions of the same blurred and warped component. This yields the following degradation model:

$$(\forall r \in \{1,\ldots,R\})(\forall b \in \{1,\ldots,B\}) \quad z_{B(r-1)+b} = \mathrm{D}_b\mathrm{T}\mathrm{W}_r\, \overline{x}_r + \varepsilon_{B(r-1)+b} \tag{5.4}$$

where $\varepsilon_{B(r-1)+b} \sim \mathcal{N}(0, \sigma^2\, \mathrm{Id}_K)$.

(iii). **Spectral unmixing** [55, 86, 161]. In this scenario, z models an hyperspectral image with $K = N$, and $\overline{x}$ is a stack of R abundance maps associated to as many spectra $(\mathrm{S}_r)_{1\le r\le R} \in (\mathbb{R}^S)^R$ known in advance. The operator A has a block diagonal structure that leads to the following mixing model:

$$(\forall \ell \in \{1,\ldots,N\}) \qquad \begin{bmatrix} z_1^{(\ell)} \\ \vdots \\ z_S^{(\ell)} \end{bmatrix} = \sum_{r=1}^{R} x_r^{(\ell)}\, \mathrm{S}_r + \varepsilon^{(\ell)} \tag{5.5}$$

where $\varepsilon^{(\ell)} \sim \mathcal{N}(0, \sigma^2\, \mathrm{Id}_S)$.

5.1.2 Convex variational approach

The usual way to recover $\overline{\mathrm{x}}$ from the observations z consists of following a convex variational approach that leads to an optimization problem in the form

$$\underset{\mathrm{x}\in C}{\text{minimize}} \quad f(\mathrm{Ax},\mathrm{z}) \quad \text{s.t.} \quad h(\mathrm{x}) \leq \eta, \tag{5.6}$$

where $\eta > 0$. The data fidelity term $f(\cdot,\mathrm{z}) \in \Gamma_0\big((\mathbb{R}^K)^S\big)$ aims at insuring that the solution is close to the observations. Depending on the noise characteristics, standard choices for f are a quadratic function for an additive Gaussian noise, an $\boldsymbol{\ell}_1$-norm when a Laplacian noise is involved, and a Kullback-Leibler divergence when dealing with Poisson noise. The function $h \in \Gamma_0\big((\mathbb{R}^N)^R\big)$ allows one to impose some regularity on the solution. Some possible choices for this function have been mentioned in the introduction. Finally, C denotes a nonempty closed convex subset of $(\mathbb{R}^N)^R$ that can be used to constrain the dynamic range of the target signal (e.g., $C = ([0,255]^N)^R$ for standard natural images).

Regarding the variational formulation of the image recovery problem, one may prefer to adopt a constrained formulation, despite state-of-the-art methods often deal with the regularized version of Problem (5.6), that is

$$\underset{\mathrm{x}\in C}{\text{minimize}} \quad f(\mathrm{Ax},\mathrm{z}) + \lambda h(\mathrm{x}), \tag{5.7}$$

where $\lambda > 0$. While both approaches are equivalent for specific values of λ and η, the constrained one might be more practical, as the choice of η can be related to some physical properties of the target signal. For example, a reasonable upper bound on the TV constraint can be available for certain classes of images, since TV constitutes a geometrical attribute that exhibits a limited variance over, e.g., views of similar urban areas in satellite imaging, tomographic reconstructions of similar cross sections, fingerprint images, text images and face images [61].

One of the difficulties of constrained approaches is that a closed form of the projection onto the involved constraint set is not always available. The second contribution of this chapter is the application of the proposed epigraphical splitting to convex optimization problems involving $\boldsymbol{\ell}_{1,p}$-matrix-norm constraints.

5.1.3 Imaging spectroscopy

Spectral imagery is used in a wide range of applications, such as remote sensing [216], astronomical imaging [172], and fluorescence microscopy [219]. In these contexts, one typically distinguishes between *multispectral* (MS) and *hyperspectral* (HS) images. In general, HS images are capable to achieve a higher spectral resolution than MS images (at the cost of acquiring a few hundred bands), which results in a better spectral characterization of the objects in the scene. This has stimulated many applications in remote sensing, such as the detection and identification of the ground surface [15]. Since HS images are characterized by the fact that an entire spectrum is acquired at each point, a huge correlation exists among close spectral bands. Consequently, a large array of variational methods have been proposed to efficiently exploit the spectral-spatial regularity present in HS images. To the best of our knowledge, these methods can be roughly divided into three classes.

A first class of methods aims at extending the regularity models used in color imagery [108, 211]. To cite a few examples, the work in [240] proposed a super-resolution method based on a component-by-component TV regularization. To deal with the huge size of HS images, the authors performed the actual super-resolution on a few principal image components (obtained by means of PCA), which are then used to interpolate the secondary components. In [30], the problem of MS denoising is dealt with by considering a hybrid regularization that induces each component to be sparse in an orthonormal basis, while promoting similarities between the components through a distance function applied on wavelet coefficients. Another kind of spectral adaptivity has been proposed in [235] for HS restoration. It consists of using the multicomponent TV regularization in [28] that averages the Frobenius norms of the multicomponent gradients. The same authors have recently proposed in [45] an inpainting method based on the multicomponent NLTV regularization. The link between this method and the proposed work will be discussed in Section 5.2.

A second class of methods aims at modeling HS images as three-dimensional tensors, i.e. two spatial dimensions and one spectral dimension. First denoising attempts in this direction were pursued in [165, 168], where tensor algebra was exploited to jointly analyze the HS hypercube by considering vectorial anisotropic diffusion methods. Other strategies consider tensorial filtering methods, such as the multiway Wiener filter (see [151] for a survey on this subject).

A third class of methods is based on robust PCA [33] or low-rank and sparse matrix decomposition [117]. These methods proceed by splitting a HS image into two separate contributions: an image formed by components having similar shapes (low-rank image), and an image that highlights the differences between the components (sparse image). For example, the work in [107] proposed a convex optimization formulation for recovering an HS image from very few compressive-sensing measurements. This approach involved a penalization based on two terms: the $\boldsymbol{\ell}_*$ nuclear norm of the matrix where each column corresponds to the 2D-wavelet coefficients of a spectral band (reshaped in a vector), and the $\boldsymbol{\ell}_{1,2}$-norm of the wavelet-coefficient blocks grouped along the spectral dimension. A similar approach was followed in [94], even though the $\boldsymbol{\ell}_* + \boldsymbol{\ell}_{1,2}$-norm penalization was applied directly on the HS pixels, rather than using a sparse transformation.

The third contribution of this chapter is the specialization of the proposed ST-NLTV regularization in the recovery of spectral images. The resulting strategy is based on tensor algebra ideas, but it uses variational strategies rather than anisotropic diffusion or adaptive filtering.

5.1.4 Outline

The chapter is organized as follows. Section 5.2 describes the nonlocal structure tensor regularization and its connections to related work. Section 5.3 explains how to employ the epigraphical splitting derived in Chapter 2 for solving Problem (5.6) via primal-dual proximal methods. Section 5.4 provides an experimental validation in the context of color, MS and HS image restoration, comparing the proposed approach with state-of-the-art methods in terms of quality and efficiency. Finally, some conclusions are drawn in Section 5.5.

5.2 Structure tensor regularization

The proposed approach consists of modeling the spatial and spectral dependencies of multicomponent images by introducing a regularization grounded on the use of *matrix norms*, leading to the following penalty

$$\big(\forall \mathrm{x} \in (\mathbb{R}^N)^R\big) \qquad h(\mathrm{x}) = \sum_{\ell=1}^{N} \tau_\ell \|\mathrm{F}_\ell \mathrm{B}_\ell \mathrm{x}\|_p, \tag{5.8}$$

where $\|\cdot\|_p$ denotes the Schatten p-norm with $p \geq 1$, $(\tau_\ell)_{1\leq\ell\leq N}$ are positive weights, and

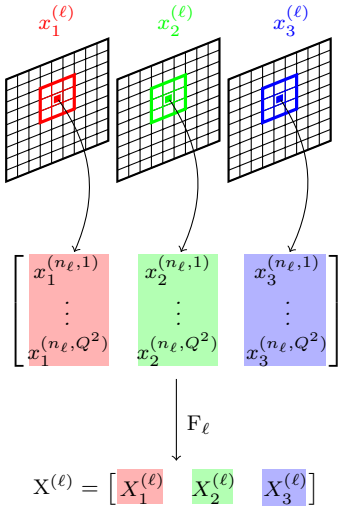

(i). **block selection**: the operator $\mathrm{B}_\ell \colon (\mathbb{R}^N)^R \to \mathbb{R}^{Q^2\times R}$ selects from each component the block of size $Q \times Q$ centered in ℓ, and rearranges them in a matrix of size $Q^2 \times R$, yielding

$$\mathrm{Y}^{(\ell)} = \begin{bmatrix} x_1^{(n_{\ell,1})} & \dots & x_R^{(n_{\ell,1})} \\ \vdots & & \vdots \\ x_1^{(n_{\ell,Q^2})} & \dots & x_R^{(n_{\ell,Q^2})} \end{bmatrix} \tag{5.9}$$

where $\mathcal{W}_\ell = \{n_{\ell,1}, \dots, n_{\ell,Q^2}\}$ denotes the set of positions located into the window around ℓ, with $Q > 1$;

(ii). **block transform**: the operator $\mathrm{F}_\ell \colon \mathbb{R}^{Q^2\times R} \to \mathbb{R}^{M_\ell \times R}$ denotes an analysis transformation that achieves a sparse representation of grouped blocks, yielding

$$\mathrm{X}^{(\ell)} = \mathrm{F}_\ell \mathrm{Y}^{(\ell)}, \tag{5.10}$$

where $M_\ell \leq Q^2$.

The resulting structure tensor regularization can be thus expressed as

$$h(\mathrm{x}) = \sum_{\ell=1}^{N} \tau_\ell \, \|\mathrm{X}^{(\ell)}\|_p. \tag{5.11}$$

By denoting the singular values of $\mathrm{X}^{(\ell)}$ ordered in decreasing order with

$$\sigma_{\mathrm{X}^{(\ell)}} = \big(\sigma^{(m)}_{\mathrm{X}^{(\ell)}}\big)_{1\leq m\leq \widetilde{M}_\ell}, \qquad \widetilde{M}_\ell = \min\{M_\ell, R\}, \tag{5.12}$$

the case $p \in [1, +\infty[$ leads to

$$h(\mathrm{x}) = \sum_{\ell=1}^{N} \tau_\ell \left(\sum_{m=1}^{\widetilde{M}_\ell} \big(\sigma^{(m)}_{\mathrm{X}^{(\ell)}}\big)^p \right)^{1/p}, \tag{5.13}$$

whereas the case $p = +\infty$ yields

$$h(\mathrm{x}) = \sum_{\ell=1}^{N} \tau_\ell \, \sigma^{(1)}_{\mathrm{X}^{(\ell)}}. \tag{5.14}$$

In the case when $p = 1$, the Schatten norm reduces to the nuclear norm, and thus the structure tensor regularization proposed in (5.10) induces a low-rank approximation of matrices $(\mathrm{X}^{(\ell)})_{1\leq\ell\leq N}$ (see [202] for a survey on singular value decomposition). In addition, the proposed regularization generalizes several state-of-the-art regularization strategies, as explained in the following.

5.2.1 ST-TV

The multicomponent TV regularization [28, 108, 235] can be retrieved by setting F_ℓ to the operator which, for each $r \in \{1,\dots,R\}$, computes the difference between $x_r^{(\ell)}$ and its horizontal/vertical nearest neighbours $(x_r^{(\ell_1)}, x_r^{(\ell_2)})$, yielding the matrix

$$\mathrm{X}_{\mathrm{TV}}^{(\ell)} = \begin{bmatrix} x_1^{(\ell)} - x_1^{(\ell_1)} & \dots & x_R^{(\ell)} - x_R^{(\ell_1)} \\ x_1^{(\ell)} - x_1^{(\ell_2)} & \dots & x_R^{(\ell)} - x_R^{(\ell_2)} \end{bmatrix} \tag{5.15}$$

with $M_\ell = 2$ (which implies $Q = 2$). A special case of ST-TV arising when $p = 2$ can be found in [235] (later referred to as *Hyperspectral-TV*). The regularization in [145] is an extension of ST-TV given by $\mathrm{X}^{(\ell)} = [\mathrm{X}_{\mathrm{TV}}^{(n)}]_{n\in\mathcal{W}_\ell}$, yielding a matrix of size $2 \times RQ^2$ (see below (5.9) for the definition of $\mathcal{W}_\ell$). Finally, note that the regularization used in [240] differs from ST-TV, as it amounts to summing up the smoothed TV [7] evaluated separately over each component.

5.2.2 ST-NLTV

The NLTV regularization [105] can be extended to multicomponent images by setting F_ℓ to the operator that, for each $r \in \{1,\dots,R\}$, computes the weighted difference between $x_r^{(\ell)}$ and some other pixel values. This results in the matrix

$$\mathrm{X}_{\mathrm{NLTV}}^{(\ell)} = \left[\omega_{\ell,n}(x_r^{(\ell)} - x_r^{(n)})\right]_{n\in\mathcal{N}_\ell, 1\le r\le R}, \tag{5.16}$$

where $\mathcal{N}_\ell \subset \mathcal{W}_\ell \setminus \{\ell\}$ denotes the non-local support of the neighborhood of ℓ, and M_ℓ corresponds to the size of this support. The regularization in [45] (later referred to as *Multichannel-NLTV*) appears as a special case of the proposed ST-NLTV arising when $p = 2$ and the local window is fully used ($M_\ell = Q^2$).

For every $\ell \in \{1,\dots,N\}$ and $n \in \mathcal{N}_\ell$, the weight $\omega_{\ell,n} > 0$ depends on the similarity between patches built around the pixels ℓ and n of the image to be recovered. Since the degradation process in (5.1) may involve some missing data, a two-step approach is adopted in order to estimate these weights. In the first step, the ST-TV regularization is used to obtain an estimate $\widetilde{\mathrm{x}}$ of the target image. This estimate subsequently serves in the second step to compute the weights through a *self-similarity* measure as follows:

$$\omega_{\ell,n} = \widetilde{\omega}_\ell \exp\left(-\delta^{-2}\,\|\mathrm{L}_\ell\widetilde{\mathrm{x}} - \mathrm{L}_n\widetilde{\mathrm{x}}\|_2^2\right), \tag{5.17}$$

where $\delta > 0$, L_ℓ (resp. L_n) selects a $\widetilde{Q} \times \widetilde{Q} \times R$ patch centered at position ℓ (resp. n) after a linear processing depending on the position ℓ (resp. n), and the constant $\widetilde{\omega}_\ell > 0$ is set so as to normalize the weights (i.e. $\sum_{n\in\mathcal{N}_\ell}\omega_{\ell,n} = 1$). Note that the linear processing is applied to improve the reliability of the self-similarity measure, and thus to insure a better image recovery performance. In the simplest case, it consists of point-wise multiplying the selected patches by a bivariate Gaussian function [32]. A more sophisticated processing may involve a convolution with a set of low-pass Gaussian filters whose variances increase as the spatial distance from the patch center grow [99]. For every $\ell \in \{1,\dots,N\}$, the neighborhood $\mathcal{N}_\ell$ is built according to the procedure described in [104], which allows one to limit the size M_ℓ of the neighborhood to a fixed value $\overline{M}$ (in all the experiments, we set $\tau_\ell \equiv 1$, $Q = 11$, $\widetilde{Q} = 5$, $\delta = 35$ and $\overline{M} = 14$).

5.3 Optimization method

Within the proposed framework, Problem (5.6) can be reformulated as follows

$$\underset{\mathrm{x}\in C}{\text{minimize}} \quad f(\mathrm{Ax},\mathrm{z}) \quad \text{s.t.} \quad \Phi\,\mathrm{x}\in D, \tag{5.18}$$

where $M=M_1+\cdots+M_N$, $\Phi\colon (\mathbb{R}^N)^R\mapsto\mathbb{R}^{M\times R}$ is the linear operator such that

$$\big(\forall \mathrm{x}\in(\mathbb{R}^N)^R\big) \qquad \Phi\mathrm{x}=\mathrm{X}=\begin{bmatrix}\mathrm{F}_1\mathrm{B}_1\mathrm{x}\\ \cdots\\ \mathrm{F}_N\mathrm{B}_N\mathrm{x}\end{bmatrix}\begin{matrix}\}\ \mathrm{X}^{(1)}\\ \\ \}\ \mathrm{X}^{(N)}\end{matrix} \tag{5.19}$$

and D is the closed convex set defined as

$$D=\big\{\mathrm{X}\in\mathbb{R}^{M\times R}\ \big|\ \sum_{\ell=1}^{N}\tau_\ell\|\mathrm{X}^{(\ell)}\|_p\leq\eta\big\}. \tag{5.20}$$

The standard way of solving Problem (5.18) with M+LFBF leads to Algorithm 5.1, whose iterations involve the projection onto D. While the latter can be computed through specific numerical procedures for $p\in\{1,2,+\infty\}$ [71, 95, 188, 217, 223], a more efficient approach consists of resorting to the proposed epigraphical splitting. The latter amounts to reformulating Problem (5.18) as

$$\underset{(\mathrm{x},\zeta)\in C\times W}{\text{minimize}} \quad f(\mathrm{Ax},\mathrm{z}) \quad \text{s.t.} \quad (\Phi\,\mathrm{x},\,\zeta)\in E, \tag{5.21}$$

where the convex sets E and W are defined in Sections 5.3.1 and 5.3.2 according to the type of Shatten p-norms involved in D. The iterations of M+LFBF associated with Problem (5.21) are summarized in Algorithm 5.2.

5.3.1 Nuclear norm

Since the case $p=1$ yields

$$\mathrm{X}\in D \quad\Leftrightarrow\quad \sum_{\ell=1}^{N}\sum_{m=1}^{\widetilde{M}_\ell}\tau_\ell\left|\sigma^{(m)}_{\mathrm{X}^{(\ell)}}\right|\leq\eta, \tag{5.22}$$

the constraint D can be decomposed by introducing an auxiliary vector $\zeta=(\zeta^{(\ell,m)})_{1\leq\ell\leq L,1\leq m\leq\widetilde{M}_\ell}$ such that

$$\begin{cases}(\forall\ell\in\{1,\ldots,N\})(\forall m\in\{1,\ldots,\widetilde{M}_\ell\}) & \left|\sigma^{(m)}_{\mathrm{X}^{(\ell)}}\right|\leq\zeta^{(\ell,m)},\\ \displaystyle\sum_{\ell=1}^{N}\sum_{m=1}^{\widetilde{M}_\ell}\tau_\ell\,\zeta^{(\ell,m)}\leq\eta. & \end{cases} \tag{5.23}$$

Consequently, the sets E and W in Problem (5.21) are defined as follows:

$$E=\big\{(\mathrm{X},\zeta)\in\mathbb{R}^{M\times R}\times\mathbb{R}^{\widetilde{M}}\mid(\forall\ell\in\{1,\ldots,N\})\\ (\forall m\in\{1,\ldots,\widetilde{M}_\ell\})\quad(\sigma^{(m)}_{\mathrm{X}^{(\ell)}},\zeta^{(\ell,m)})\in\mathrm{epi}\,|\cdot|\,\big\}, \tag{5.24}$$

$$W=\big\{\zeta\in\mathbb{R}^{\widetilde{M}}\ \big|\ \sum_{\ell=1}^{N}\sum_{m=1}^{\widetilde{M}_\ell}\tau_\ell\,\zeta^{(\ell,m)}\leq\eta\big\}, \tag{5.25}$$

and the epigraphical projection is given in Section 5.3.3.

5.3.2 Frobenius and spectral norms

Since the case $p > 1$ leads to

$$\mathrm{X} \in D \quad \Leftrightarrow \quad \sum_{\ell=1}^{N} \tau_\ell \, \|\sigma_{\mathrm{X}^{(\ell)}}\|_p \leq \eta, \tag{5.26}$$

the constraint D can be decomposed by introducing an auxiliary vector $\zeta = (\zeta^{(\ell)})_{1\leq\ell\leq N}$ such that

$$\begin{cases} (\forall \ell \in \{1,\dots,N\}) \quad \|\sigma_{\mathrm{X}^{(\ell)}}\|_p \leq \zeta^{(\ell)}, \\ \displaystyle\sum_{\ell=1}^{N} \tau_\ell \, \zeta^{(\ell)} \leq \eta. \end{cases} \tag{5.27}$$

Consequently, the sets E and W in Problem (5.21) are defined as follows:

$$\begin{aligned} E &= \big\{(\mathrm{X},\zeta) \in \mathbb{R}^{M\times R} \times \mathbb{R}^N \;\big|\; (\forall \ell \in \{1,\dots,N\}) \quad (\sigma_{\mathrm{X}^{(\ell)}}, \zeta^{(\ell)}) \in \operatorname{epi} \|\cdot\|_p \big\}, \\ W &= \big\{\zeta \in \mathbb{R}^N \;\big|\; \sum_{\ell=1}^{N} \tau_\ell \, \zeta^{(\ell)} \leq \eta \big\}, \end{aligned} \tag{5.28}$$

and the epigraphical projection is given in Section 5.3.3.

5.3.3 Epigraphical projection

The advantage of resorting to the epigraphical splitting is that the projection onto E admits a closed-form expression in the case $p \in \{1, 2, +\infty\}$, as summarized in the following proposition (which is straightforwardly proved).

Proposition 5.3.1. *For every $\ell \in \{1,\dots,N\}$, let*

$$\mathrm{X}^{(\ell)} = \mathrm{U}^{(\ell)} \mathrm{S}^{(\ell)} \big(\mathrm{V}^{(\ell)}\big)^\top \tag{5.29}$$

be the Singular Value Decomposition of $\mathrm{X}^{(\ell)} \in \mathbb{R}^{M_\ell \times R}$, where

- $(\mathrm{U}^{(\ell)})^\top \mathrm{U}^{(\ell)} = \mathrm{Id}_{\widetilde{M}_\ell}$,
- $\big(\mathrm{V}^{(\ell)}\big)^\top \mathrm{V}^{(\ell)} = \mathrm{Id}_{\widetilde{M}_\ell}$,
- $\mathrm{S}^{(\ell)} = \mathrm{Diag}(\mathsf{s}^{(\ell)})$, *with* $\mathsf{s}^{(\ell)} = (\sigma^{(m)}_{\mathrm{X}^{(\ell)}})_{1\leq m\leq \widetilde{M}_\ell}$.

The SVD can be avoided if p=2, as the Frobenius norm is the ℓ_2-norm of the vector formed by all matrix elements [202].

Then,

$$P_E(\mathrm{X},\zeta) = \Big(\mathrm{U}^{(\ell)} \mathrm{T}^{(\ell)} \big(\mathrm{V}^{(\ell)}\big)^\top, \; \theta^{(\ell)}\Big)_{1\leq\ell\leq N}, \tag{5.30}$$

where $\mathrm{T}^{(\ell)} = \mathrm{Diag}(\mathsf{t}^{(\ell)})$ and

$$(\mathsf{t}^{(\ell)}, \; \theta^{(\ell)}) = \begin{cases} \big[P_{\operatorname{epi}|\cdot|}(\mathsf{s}^{(\ell,m)}, \zeta^{(\ell,m)})\big]_{1\leq m\leq \widetilde{M}_\ell}, & \text{if } p = 1, \\ P_{\operatorname{epi}\|\cdot\|_p}(\mathsf{s}^{(\ell)}, \zeta^{(\ell)}), & \text{if } p > 1, \end{cases} \tag{5.31}$$

with the above epigraphical projections given in Propositions 2.3.2, 2.3.5, and 2.3.7.

Algorithm 5.1 *M+LFBF for solving Problem* (5.18)

INITIALIZATION

$$\left\lfloor \begin{array}{l} \text{choose } (\mathrm{x}^{[0]}, \mathrm{Y}_1^{[0]}, \mathrm{y}_2^{[0]}) \in (\mathbb{R}^N)^R \times \mathbb{R}^{M\times R} \times (\mathbb{R}^K)^S \\ \text{set } \gamma \in \left]0, \left(\|\mathrm{A}\|^2 + \|\Phi\|^2\right)^{-1/2}\right[\end{array} \right.$$

FOR $t = 0, 1, \ldots$

$$\left\lfloor \begin{array}{l} \mathrm{p}^{[t]} = P_C\Big(\mathrm{x}^{[t]} - \gamma(\Phi^\top \mathrm{Y}_1^{[t]} + \mathrm{A}^\top \mathrm{y}_2^{[t]})\Big) \\ \widehat{\mathrm{Y}}_1^{[t]} = \mathrm{Y}_1^{[t]} + \gamma\Phi\mathrm{x}^{[t]} \\ \widetilde{\mathrm{Y}}_1^{[t]} = \widehat{\mathrm{Y}}_1^{[t]} - \gamma P_D\Big(\widehat{\mathrm{Y}}_1^{[t]}/\gamma\Big) \\ \mathrm{Y}_1^{[t+1]} = \widetilde{\mathrm{Y}}_1^{[t]} + \gamma\Phi\big(\mathrm{p}^{[t]} - \mathrm{x}^{[t]}\big) \\ \widehat{\mathrm{y}}_2^{[t]} = \mathrm{y}_2^{[t]} + \gamma\mathrm{A}\mathrm{x}^{[t]} \\ \widetilde{\mathrm{y}}_2^{[t]} = \widehat{\mathrm{y}}_2^{[t]} - \gamma\,\mathrm{prox}_{\gamma^{-1}f(\cdot;\mathrm{z})}\Big(\widehat{\mathrm{y}}_2^{[t]}/\gamma\Big) \\ \mathrm{y}_2^{[t+1]} = \widetilde{\mathrm{y}}_2^{[t]} + \gamma\mathrm{A}\big(\mathrm{p}^{[t]} - \mathrm{x}^{[t]}\big) \\ \mathrm{x}^{[t+1]} = \mathrm{p}^{[t]} - \gamma\Phi^\top\big(\widetilde{\mathrm{Y}}_1^{[t]} - \mathrm{Y}_1^{[t]}\big) - \gamma\mathrm{A}^\top\big(\widetilde{\mathrm{y}}_2^{[t]} - \mathrm{y}_2^{[t]}\big) \end{array} \right.$$

Algorithm 5.2 *M+LFBF for solving Problem* (5.21)

INITIALIZATION

$$\left\lfloor \begin{array}{l} \text{choose } (\mathrm{x}^{[0]}, \mathrm{Y}_1^{[0]}, \mathrm{y}_2^{[0]}) \in (\mathbb{R}^N)^R \times \mathbb{R}^{M\times R} \times (\mathbb{R}^K)^S \\ \text{choose } (\zeta^{[0]}, \nu_1^{[0]}) \in W \times W \\ \text{set } \gamma \in \left]0, \left(\|\mathrm{A}\|^2 + \max\{\|\Phi\|^2, 1\}\right)^{-1/2}\right[\end{array} \right.$$

FOR $t = 0, 1, \ldots$

$$\left\lfloor \begin{array}{l} \mathrm{p}^{[t]} = P_C\Big(\mathrm{x}^{[t]} - \gamma(\Phi^\top \mathrm{Y}_1^{[t]} + \mathrm{A}^\top \mathrm{y}_2^{[t]})\Big) \\ \rho^{[t]} = P_W\big(\zeta^{[t]} - \gamma\nu_1^{[t]}\big) \\ \widehat{\mathrm{Y}}_1^{[t]} = \mathrm{Y}_1^{[t]} + \gamma\Phi\mathrm{x}^{[t]} \\ \widehat{\nu}_1^{[t]} = \nu_1^{[t]} + \gamma\zeta^{[t]} \\ \Big(\widetilde{\mathrm{Y}}_1^{[t]}, \widetilde{\nu}_1^{[t]}\Big) = \Big(\widehat{\mathrm{Y}}_1^{[t]}, \widehat{\nu}_1^{[t]}\Big) - \gamma P_E\Big(\widehat{\mathrm{Y}}_1^{[t]}/\gamma, \widehat{\nu}_1^{[t]}/\gamma\Big) \\ \mathrm{Y}_1^{[t+1]} = \widetilde{\mathrm{Y}}_1^{[t]} + \gamma\Phi\big(\mathrm{p}^{[t]} - \mathrm{x}^{[t]}\big) \\ \nu_1^{[t+1]} = \widetilde{\nu}_1^{[t]} + \gamma\big(\rho^{[t]} - \zeta^{[t]}\big) \\ \widehat{\mathrm{y}}_2^{[t]} = \mathrm{y}_2^{[t]} + \gamma\mathrm{A}\mathrm{x}^{[t]} \\ \widetilde{\mathrm{y}}_2^{[t]} = \widehat{\mathrm{y}}_2^{[t]} - \gamma\,\mathrm{prox}_{\gamma^{-1}f(\cdot;\mathrm{z})}\Big(\widehat{\mathrm{y}}_2^{[t]}/\gamma\Big) \\ \mathrm{y}_2^{[t+1]} = \widetilde{\mathrm{y}}_2^{[t]} + \gamma\mathrm{A}\big(\mathrm{p}^{[t]} - \mathrm{x}^{[t]}\big) \\ \mathrm{x}^{[t+1]} = \mathrm{p}^{[t]} - \gamma\Phi^\top\big(\widetilde{\mathrm{Y}}_1^{[t]} - \mathrm{Y}_1^{[t]}\big) - \gamma\mathrm{A}^\top\big(\widetilde{\mathrm{y}}_2^{[t]} - \mathrm{y}_2^{[t]}\big) \\ \zeta^{[t+1]} = \rho^{[t]} - \gamma\big(\widetilde{\nu}_1^{[t]} - \nu_1^{[t]}\big) \end{array} \right.$$

5.3.4 Approach based on ADMM

An alternative approach to deal with Problem (5.18) consists of employing the Alternating Direction Method of Multipliers (ADMM) [91] or one of its parallel versions [2, 64, 65, 90, 180, 198], sometimes referred to as the Simultaneous Direction Method of Multipliers (SDMM). Although these algorithms are appealing, they require to invert the operator $\mathrm{Id} + \Phi^\top \Phi + \mathrm{A}^\top \mathrm{A}$, which makes them less attractive than primal-dual algorithms for solving Problem (5.18). Indeed, this matrix is not diagonalizable in the DFT domain (due to the form of Φ), which rules out efficient inversion techniques, such as those employed in [1, 2, 187]. When Φ denotes the NLTV operator defined in (5.16), this issue can be circumvented [30, 182] by decomposing it as $\Phi = \Omega G$, where $\mathrm{G}\colon (\mathbb{R}^N)^R \to \mathbb{R}^{M \times N(Q^2-1)}$ is a concatenation of discrete difference operators, and $\Omega \in \mathbb{R}^{M \times N(Q^2-1)}$ is a weighted block-selection operator. So doing, Problem (5.18) can be equivalently reformulated by introducing an auxiliary variable $\xi = \Phi \mathrm{x} \in \mathbb{R}^{M \times R}$, yielding

$$\underset{(\mathrm{x},\xi) \in C \times D}{\text{minimize}} \quad f(\mathrm{Ax}, \mathrm{z}) \qquad \text{s.t.} \qquad (\mathrm{Gx}, \xi) \in V, \tag{5.32}$$

where $V = \big\{(\mathrm{X}, \xi) \in \mathbb{R}^{N(Q^2-1) \times R} \times \mathbb{R}^{M \times R} \;\big|\; \Omega \mathrm{X} = \xi\big\}$. The iterations of SDMM associated to Problem (5.32) are illustrated in Algorithm 5.3, where the inversion of matrix $H = \mathrm{Id} + \mathrm{G}^\top \mathrm{G} + \mathrm{A}^\top \mathrm{A}$ is performed in the DFT domain.

It is worth emphasizing that SDMM still requires to compute the projection onto D, which may be done by either resorting to specific numerical solutions [71, 95, 188, 217, 223] or employing the proposed epigraphical splitting. However, according to our simulations, both approaches are slower than Algorithm 5.2.

Algorithm 5.3 *SDMM for solving Problem* (5.32)

Initialization

$\mathrm{y}_1^{[0]} \in (\mathbb{R}^N)^R, \mathrm{Y}_2^{[0]} \in \mathbb{R}^{M \times R}, \mathrm{y}_3^{[0]} \in (\mathbb{R}^K)^S$

$\overline{\mathrm{y}}_1^{[0]} \in (\mathbb{R}^N)^R, \overline{\mathrm{Y}}_2^{[0]} \in \mathbb{R}^{M \times R}, \overline{\mathrm{y}}_3^{[0]} \in (\mathbb{R}^K)^S$

$\chi_1^{[0]} \in \mathbb{R}^{M \times R}, \chi_2^{[0]} \in \mathbb{R}^{M \times R}$

$\overline{\chi}_1^{[0]} \in \mathbb{R}^{M \times R}, \overline{\chi}_2^{[0]} \in \mathbb{R}^{M \times R}$

$\mathrm{H} = \mathrm{Id} + \mathrm{G}^\top \mathrm{G} + \mathrm{A}^\top \mathrm{A}$

For $t = 0, 1, \ldots$

$\gamma_t \in]0, +\infty[$

$\mathrm{x}^{[t]} = \mathrm{H}^{-1} \left[\mathrm{y}_1^{[t]} - \overline{\mathrm{y}}_1^{[t]} + \mathrm{G}^\top (\mathrm{Y}_2^{[t]} - \overline{\mathrm{Y}}_2^{[t]}) + \mathrm{A}^\top (\mathrm{y}_3^{[t]} - \overline{\mathrm{y}}_3^{[t]}) \right]$

$\xi^{[t]} = \frac{1}{2} \left(\chi_1^{[t]} - \overline{\chi}_1^{[t]} \right) + \frac{1}{2} \left(\chi_2^{[t]} - \overline{\chi}_2^{[t]} \right)$

$\mathrm{y}_1^{[t+1]} = P_C \left(\mathrm{x}^{[t]} + \overline{\mathrm{y}}_1^{[t]} \right)$

$\chi_1^{[t+1]} = P_D \left(\xi^{[t]} + \overline{\chi}_1^{[t]} \right)$

$\left(\mathrm{Y}_2^{[t+1]}, \chi_2^{[t+1]} \right) = P_V (\mathrm{Gx}^{[t]} + \overline{\mathrm{Y}}_2^{[t]}, \xi^{[t]} + \overline{\chi}_2^{[t]})$

$\mathrm{y}_3^{[t+1]} = \mathrm{prox}_{\gamma_t f} \left(\mathrm{Ax}^{[t]} + \overline{\mathrm{y}}_3^{[t]} \right)$

$\overline{\mathrm{y}}_1^{[t+1]} = \overline{\mathrm{y}}_1^{[t]} + \mathrm{x}^{[t]} - \mathrm{y}_1^{[t+1]}$

$\overline{\chi}_1^{[t+1]} = \overline{\chi}_1^{[t]} + \xi^{[t]} - \chi_1^{[t+1]}$

$\left(\overline{\mathrm{Y}}_2^{[t+1]}, \overline{\chi}_2^{[t+1]} \right) = \left(\overline{\mathrm{Y}}_2^{[t]} + \mathrm{Gx}^{[t]} - \mathrm{Y}_2^{[t+1]}, \overline{\chi}_2^{[t]} + \xi^{[t]} - \chi_2^{[t+1]} \right)$

$\overline{\mathrm{y}}_3^{[t+1]} = \overline{\mathrm{y}}_3^{[t]} + \mathrm{Ax}^{[t]} - \mathrm{y}_3^{[t+1]}$

5.4 Numerical results

A numerical analysis of the proposed approach is now conducted. In Sections 5.4.1 and 5.4.2, the visual impact of the proposed regularization is evaluated on color and spectral images. In Section 5.4.3, the execution time of Algorithm 5.2 is assessed w.r.t. some alternatives. In all the experiments, the pixels of each component in the original image are normalized in $[0, 255]$, hence the dynamic range constraint set C imposes that the pixel values belong to such an interval. Moreover, the noise is assumed to be additive white Gaussian, thus the fidelity term related to its anti-log likelihood is $f = \|\mathrm{A} \cdot -\mathrm{z}\|_2^2$.

5.4.1 Color photography

The first experiment is focused on color imaging (i.e., $R = S = 3$). The noisy observations are obtained with the degradation model in (5.3), where the measurement operator D_r denotes a decimated convolution. While it is common for color imaging to work in a luminance-chrominance space [70, 72], or a perceptually-uniform space [108], the random decimation prevents us from taking this approach, as pixels having missing colors cannot be correctly projected onto a different color space. The tests are thus conducted in the RGB space.

Figure 5.1 collects the images reconstructed by using, for every $p \in \{1, 2, +\infty\}$, the proposed $\boldsymbol{\ell}_p$-ST-NLTV, the standard $\boldsymbol{\ell}_p$-ST-TV [28, 108], and the more classical *channel-by-channel* (CC) regularization [6, 51], defined as

$$h_{\mathrm{cc}}(\mathrm{x}) = \sum_{r=1}^{R} \sum_{\ell=1}^{N} \|X_r^{(\ell)}\|_p, \tag{5.33}$$

where $X_r^{(\ell)}$ denotes the r-th column vector of matrix $\mathrm{X}^{(\ell)}$ in (5.10). For all the considered constraints, the bounds η were hand-tuned in order to obtain the best SNR values. The results indicate that $\boldsymbol{\ell}_1$-ST-NLTV is the most effective regularization, because it combines the advantages of both ST and NLTV, namely a better preservation of details and a reduction of color smearing.

5.4.2 Imaging spectroscopy

In this section, the proposed $\boldsymbol{\ell}_1$-ST-NLTV is compared with implementations of two state-of-the-art methods in spectral imagery: Hyperspectral TV (H-TV) [235] (see Section 5.2.1), and Multichannel NLTV (M-NLTV) [45] (see Section 5.2.2). To this end, two scenarios are addressed by using the degradation model in (5.3): a compressive-sensing scenario in which the measurement operator $(\mathrm{D}_r)_{1 \le r \le R}$ is a random decimation, and a restoration scenario in which $(\mathrm{D}_r)_{1 \le r \le R}$ is a decimated convolution. For reproducibility purposes, several publicly available multispectral and hyperspectral images were used (see online at `engineering.purdue.edu/~biehl/MultiSpec/hyperspectral.html`). The SNR index is used to give a quantitative assessment of the results obtained from the simulated experiments, reporting both the SNR computed over all the image, and the average of SNRs evaluated component-by-component (M-SNR).

Extensive tests have been carried out on several images of different sizes. Tables 5.1 and 5.2 collect the SNR and M-SNR indexes obtained with H-TV, M-

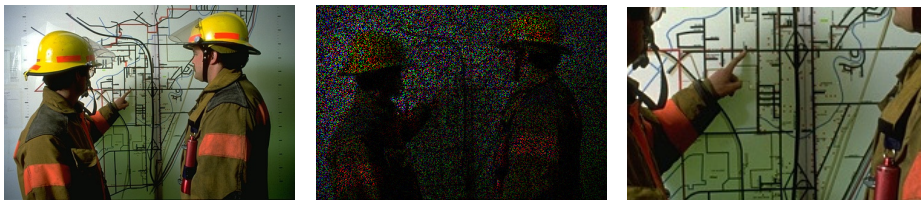

(a) Original. (b) Noisy. (c) Zoom.

(d) ℓ_1-CC-TV [6]: 16.15 dB (e) ℓ_2-CC-TV [6]: 16.32 dB (f) ℓ_∞-CC-TV [6]: 16.05 dB

(g) ℓ_1-CC-NLTV[51]: 16.87 dB (h) ℓ_2-CC-NLTV[51]: 17.20 dB (i) ℓ_∞-CC-NLTV[51]: 17.22 dB

(j) ℓ_1-ST-TV: 17.08 dB (k) ℓ_2-ST-TV [28]: 16.84 dB (l) ℓ_∞-ST-TV [108]: 16.43 dB

(m) ℓ_1-ST-NLTV: **18.20 dB** (n) ℓ_2-ST-NLTV: 17.46 dB (o) ℓ_∞-ST-NLTV: 16.67 dB

Figure 5.1 *Visual comparison of a color image reconstructed with various regularization constraints. Degradation: additive zero-mean white Gaussian noise with std. deviation equal to* 10*, uniform blur of size* 3×3*, and* 80% *of decimation (*$N = 154401$*,* $R = S = 3$ *and* $K = 30880$*).*

NLTV, ℓ_1-ST-TV, and ℓ_1-ST-NLTV, for the two degradation scenarios mentioned above. The hyper-parameter for each method (the bound η for the ST constraint in our algorithm) was hand-tuned in order to achieve the best SNR values. The best results are highlighted in bold. Moreover, a component-by-component comparison of two hyperspectral images is made in Figure 5.3, while a visual comparison of a component from the image *hydice* is displayed in Figure 5.2.

The aforementioned results demonstrate the interest of combining the non-locality principle with measures based on the structure tensor. Indeed, ℓ_1-ST-NLTV proves to be the most effective regularization with gains in SNR (up to 1.4 dB) with respect to M-NLTV, which in turn is comparable with ℓ_1-ST-TV. The better performance of ℓ_1-ST-NLTV seems to be related to its ability to better preserve edges and thin structures, while preventing component smearing.

Table 5.1 *SNR – M-SNR indexes resulting from the reconstruction of images degraded by a Gaussian noise with std. deviation equal to* 5*, and* 90% *of decimation.*

image	size	H-TV [235]	ℓ_1-ST-TV	M-NLTV [45]	ℓ_1-ST-NLTV
Hydice	$256 \times 256 \times 191$	10.65 – 09.87	11.93 – 11.16	11.57 – 10.76	**12.98 – 12.11**
Indian Pine	$145 \times 145 \times 200$	17.31 – 17.00	18.46 – 18.24	17.62 – 17.34	**19.53 – 19.49**
Little River	$512 \times 512 \times 7$	17.81 – 18.20	18.49 – 18.83	18.46 – 18.90	**19.88 – 20.18**
Mississippi	$512 \times 512 \times 7$	18.27 – 18.07	18.60 – 18.37	18.94 – 18.59	**19.56 – 19.28**
Montana	$512 \times 512 \times 7$	22.49 – 20.97	22.68 – 21.15	22.85 – 21.29	**23.31 – 21.76**
Rio	$512 \times 512 \times 7$	16.48 – 15.29	16.65 – 15.48	16.82 – 15.64	**17.20 – 16.05**
Paris	$512 \times 512 \times 7$	14.85 – 14.31	14.94 – 14.39	15.05 – 14.53	**15.36 – 14.82**

Table 5.2 *SNR – M-SNR indexes resulting from the reconstruction of images degraded by a Gaussian noise with std. deviation equal to* 5*,* 5×5 *uniform blur, and* 70% *of decimation.*

image	size	H-TV [235]	ℓ_1-ST-TV	M-NLTV [45]	ℓ_1-ST-NLTV
Hydice	$256 \times 256 \times 191$	13.76 – 12.90	14.30 – 13.50	13.84 – 12.98	**14.84 – 14.08**
Indian Pine	$145 \times 145 \times 200$	19.80 – 19.65	20.22 – 20.13	19.73 – 19.57	**20.43 – 20.41**
Little River	$512 \times 512 \times 7$	21.35 – 21.88	21.62 – 22.01	21.31 – 22.00	**21.99 – 22.49**
Mississippi	$512 \times 512 \times 7$	21.12 – 20.29	21.21 – 20.27	21.41 – 20.52	**21.65 – 20.83**
Montana	$512 \times 512 \times 7$	24.80 – 23.37	24.82 – 23.31	24.96 – 23.53	**25.18 – 23.72**
Rio	$512 \times 512 \times 7$	18.62 – 17.50	18.57 – 17.48	18.57 – 17.60	**18.87 – 17.80**
Paris	$512 \times 512 \times 7$	16.68 – 16.55	16.80 – 16.53	16.73 – 16.60	**17.05 – 16.81**

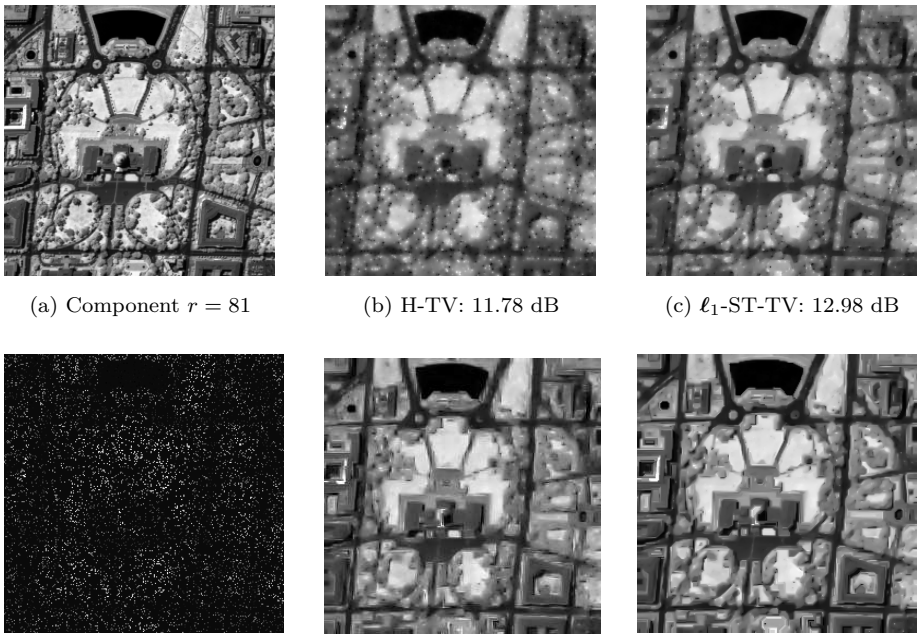

(a) Component $r = 81$ (b) H-TV: 11.78 dB (c) ℓ_1-ST-TV: 12.98 dB

(d) Noisy. (e) M-NLTV: 12.76 dB (f) ℓ_1-ST-NLTV: **14.36 dB**

Figure 5.2 *Comparison of the HS image* hydice *reconstructed with H-TV [235], ℓ_1-ST-TV, M-NLTV [45] and ℓ_1-ST-NLTV. Degradation: compressive sensing scenario involving an additive zero-mean white Gaussian noise with std. deviation* 5 *and* 90% *of decimation (*$N = 65536$, $R = 191$, $K = 6553$ *and* $S = 191$*).*

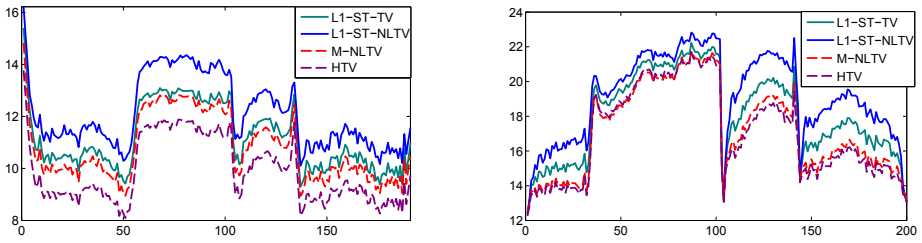

(a) SNR (dB) vs component index (image: hydice).

(b) SNR (dB) vs component index (image: indian pine).

Figure 5.3 *Quantitative comparison of two hyperspectral images reconstructed with H-TV [235], ℓ_1-ST-TV, M-NLTV [45] and ℓ_1-ST-NLTV. Degradation: compressive sensing scenario involving an additive zero-mean white Gaussian noise with std. deviation* 5 *and* 90% *of decimation.*

5.4.3 Comparison with SDMM

To complete our analysis, the execution time of Algorithm 5.2 is compared with three alternative solutions:

- Algorithm 5.1, which is a specialization of M+LFBF to Problem (5.18) where the projection onto D is computed via the procedure in [217];
- Algorithm 5.3, which is a specialization of SDMM to Problem (5.32) where the projection onto D is computed via the procedure in [217];
- the epigraphical version of Algorithm 5.3, namely SDMM applied to Problem (5.32) after that the constraint D is replaced by E and W.

It is worth emphasizing that all the above algorithms solve *exactly* Problem (5.18), hence they produce equivalent results. Our objective here is to empirically demonstrate that the epigraphical splitting technique and primal-dual proximal algorithms constitute a competitive choice for the problem at hand.

The results refer to the image *indian pine*, since a similar behavior was observed for other images. The stopping criterion is set to $\|x^{[i+1]} - x^{[i]}\| \leq 10^{-5}\|x^{[i]}\|$. The basic structure of the aforementioned algorithms was developed in Matlab, with the most "complex" operations (such as the non-local gradient and projection computations) being implemented in C using mex files. In order to compute the projection onto D, the $\boldsymbol{\ell}_1$-ball projector in [217, Algorithm 2] was used, as it avoids the expensive sorting operation (a review of several $\boldsymbol{\ell}_1$-ball projectors can be found in [71]). The codes were executed in Matlab R2011b with an Intel Xeon CPU at 2.80 GHz and 8 GB of RAM.

Figure 5.4 shows the relative error $\|x^{[i]} - x^{[\infty]}\|/\|x^{[\infty]}\|$ as a function of the computational time, where $x^{[\infty]}$ denotes the solution computed with a stopping criterion of 10^{-5} (up to 10^6 iterations). These plots indicate that the epigraphical approach yields a faster convergence than the direct one for both SDMM and M+LFBF, the latter being much faster than the former. This can be explained by the computational cost of the subiterations required by the direct projection onto the $\boldsymbol{\ell}_1$-ball. Note that these conclusions extend to all images in the dataset.

The results in Figure 5.4 refer to the bound η that achieves the best SNR indices. In practice, the optimal bound may not be known precisely, although a reasonable estimate may be available for certain classes of images based on statistics of databases [61]. While it is out of the scope of this paper to investigate an optimal strategy to set this bound, it is important to evaluate the impact of its choice on our method performance. Tables 5.3 and 5.4 compare the epigraphical approach with the direct computation of the projections (via standard iterative solutions) for different choices of η. For better readability, the values of η are expressed as a multiplicative factor of the ST-TV and ST-NLTV semi-norms of the original image. The execution times indicate that the epigraphical approach yields a faster convergence than the direct approach for SDMM and M+LFBF. Moreover, the numerical results show that errors within $\pm 5\%$ from the optimal value for η lead to SNR variations within 1.2%.

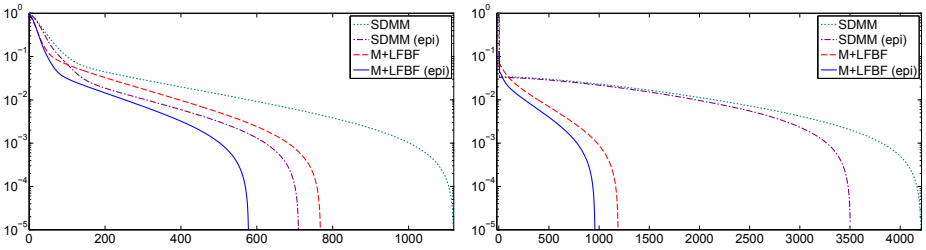

(a) ℓ_1-ST-TV. (b) ℓ_1-ST-NLTV.

Figure 5.4 *Comparison between epigraphical and direct methods:* $\frac{\|x^{[i]}-x^{[\infty]}\|}{\|x^{[\infty]}\|}$ *vs time (Degradation: std. deviation* $= 5$*, decimation* $= 90\%$*).*

Table 5.3 *Results for the ℓ_1-ST-TV constraint and some values of η. Degradation: std. deviation $= 5$, decimation $= 90\%$ ("speed up" is the ratio between "direct" and "epigraphical" times)*

η	SNR (dB) – M-SNR (dB)	SDMM					M+LFBF				
		direct		epigraphical		speed up	direct		epigraphical		speed up
		# iter.	sec.	# iter.	sec.		# iter.	sec.	# iter.	sec.	
0.35	18.41 – 18.19	547	767.51	471	466.80	1.64	466	471.95	389	339.24	1.39
0.40	**18.46 – 18.24**	**838**	**1066.24**	**698**	**701.03**	**1.52**	**733**	**735.36**	**621**	**558.37**	**1.32**
0.45	18.26 – 18.02	1000	1353.13	1000	990.76	1.37	1000	1018.58	1000	902.00	1.13

Table 5.4 *Results for the ℓ_1-ST-NLTV constraint and some values of η. Degradation: std. deviation $= 5$, decimation $= 90\%$ ("speed up" is the ratio between "direct" and "epigraphical" times)*

η	SNR (dB) – M-SNR (dB)	SDMM					M+LFBF				
		direct		epigraphical		speed up	direct		epigraphical		speed up
		# iter.	sec.	# iter.	sec.		# iter.	sec.	# iter.	sec.	
		Neighbourhood size: $Q = 3$									
0.25	19.15 – 19.05	1000	4384.38	1000	3583.57	1.23	190	494.77	190	448.22	1.11
0.30	19.39 – 19.32	1000	4414.94	1000	3417.18	1.29	243	649.31	236	534.50	1.21
0.35	19.36 – 19.28	875	4175.52	1000	3482.80	1.20	319	839.86	308	726.50	1.16
		Neighbourhood size: $Q = 5$									
0.25	19.43 – 19.38	1000	14412.86	1000	10167.34	1.42	216	977.95	212	871.80	1.12
0.30	**19.55 – 19.51**	**1000**	**14338.36**	**1000**	**10174.68**	**1.41**	**275**	**1257.71**	**268**	**1143.35**	**1.10**
0.35	19.53 – 19.49	1000	14365.92	1000	10356.73	1.39	358	1631.17	347	1424.72	1.14

5.5 Conclusions

We have proposed a new regularization for multicomponent images that is a combination of *nonlocal total variation* and *structure tensor*. The resulting image recovery problem has been formulated as a constrained convex optimization problem and solved through the epigraphical splitting and primal-dual proximal algorithms. The obtained results demonstrate the better performance of structure tensor and nonlocal gradients over a number of multispectral and hyperspectral images, although it would be interesting to consider other applications, such as the recovery of video signals or volumetric images. Our results also show that the nuclear norm has to be preferred over the Frobenius norm for hyperspectral image recovery problems. Furthermore, the experimental part indicates that the epigraphical method converges faster than the approach based on the direct computation of the projections via standard iterative solutions. In both cases, the proposed algorithm turns out to be faster than solutions based on ADMM, suggesting that primal-dual proximal algorithms constitute a good choice in practice to deal with multicomponent image recovery problems.

No amount of experimentation can ever prove me right, but a single experiment can prove me wrong.

Albert Einstein

Chapter 6
Sparse multiclass SVM

The main focus of this chapter is the learning of support vector machines through a sparse regularization and the multiclass hinge loss formulated by Crammer and Singer. To this end, we consider two approaches: one including the hinge loss as a penalty term, and the other one addressing the case when the hinge loss is enforced as a constraint. We implement the corresponding algorithms through a primal-dual proximal method and the proposed epigraphical splitting. Experiments on different types of databases, as well as comparisons with state-of-the-art methods, demonstrate the validity of our approach.

6.1 Introduction

Support vector machines (SVMs) have gained much popularity in solving classification problems with a large amount of training data or a huge number of classes [121, 126, 141, 166, 213]. While the major difficulty encountered in the large-scale SVM learning problem stems from the computational cost, this issue can be circumvented by resorting to standard Lagrangian duality techniques [75, 78]. Indeed, this approach brings in several advantages, such as the kernel trick [4], or the possibility to break the problem down into a sequence of smaller ones [19, 183]. Some works also proposed to approximate the dual problem using cutting plane approaches, in order to address scenarios with a lot of training data and thousands (or even an infinite number) of classes [126, 213].

In some applications, however, only a small number of training data is available. This is undoubtedly true in medical contexts, where the goal is to classify a patient as being "healthy", "contaminated", or "infected", but the verified cases of infected patients might be just a few. In such applications, the lack of training data may lead to the so-called *overfitting* problem, eventually yielding a prediction which is too strongly tailored to the particularities of the training set and poorly generalizes to new data.

A common solution to prevent overfitting consists of introducing a sparsity-inducing regularization, in order to perform an implicit *feature selection* that gets rid of irrelevant or noisy features. In this respect, the ℓ_1-norm and, more generally, the $\ell_{1,p}$-norm regularization have attracted much attention over the past decade [9, 88, 137, 189, 194, 214, 218, 233]. However, when a sparse regularization is introduced, the duality approach is no longer useful. Therefore, the learning of sparse SVMs leads to a nonsmooth convex optimization problem which is challenging to deal with. The main objective of this chapter is to *exactly* and *efficiently* solve the multiclass SVM learning problem involving a convex regularization and the hinge loss formulated by Crammer and Singer [78].

6.1.1 Related work

The use of sparse regularization in SVMs was firstly proposed in the context of binary classification. The idea traces back to Bradley and Mangasarian [26], who demonstrated that the ℓ_1-norm regularization can effectively perform "feature selection" by shrinking small coefficients to zero. Other forms of regularization have also been studied, such as the ℓ_0-norm [224], the ℓ_p-norm with $p > 0$ [155], the ℓ_∞-norm [245], and the combination of ℓ_0-ℓ_1 norms [156] or ℓ_1-ℓ_2 norms [221]. A different solution consists of reformulating the SVM learning problem through an indicator vector (its components being equal to 0 or 1) to model the active features [206], which leads to a combinatorial problem addressable by convex relaxation. More recently, Laporte et al. [140] proposed an accelerated algorithm for ℓ_1-regularized SVMs involving the *square hinge* loss. They also proposed a procedure for handling nonconvex regularization (using the reweighted ℓ_1-minimization scheme by [34]), showing that nonconvex penalties lead to similar prediction quality while using less features than convex ones.

Binary SVMs can be turned into multiclass classifiers by a variety of strategies, such as the *one-vs-all* approach [75, 191]. While these techniques provide a simple and powerful framework, they cannot capture the correlations between different classes, since they break a multiclass problem into multiple *independent* binary problems. Crammer and Singer [78] therefore proposed a direct formulation of multiclass SVMs by generalizing the notion of margins used in the binary case. A natural idea thus consists of equipping muticlass SVMs with sparse regularization. A simple example is the ℓ_1-regularized multiclass SVM, which can be addressed by linear programming techniques [220]. In multiclass problems, however, feature selection becomes more complex than in the binary case, since multiple discriminating functions need to be estimated, each one with its own set of important features. For this reason, mixed-norm regularization has attracted much interest due to its ability to impose group sparsity [88, 167, 175, 234].

In the context of multiclass SVMs, Zhang et al. [241] proposed to deal with the $\ell_{1,\infty}$-norm regularization by reformulating the SVM learning problem in terms of linear programming. However, they validated their method on small-size problems, indicating that the linear reformulation may be inefficient for larger-size ones. More recently, Blondel et al. [20] proposed an algorithm to handle $\ell_{1,2}$-regularized SVMs involving a smooth loss function. While their method is efficient and can handle other convex regularizations, the multiclass SVM learning is not solved rigorously, possibly leading to performance limitations.

6.1.2 Contributions

The aforementioned methods deal with sparse multiclass SVMs by either finding an approximate solution [20, 140, 206] or employing inefficient linear programming techniques [220, 241]. This chapter presents a novel approach for sparse multiclass SVMs learning based on a primal-dual proximal method [69] and the proposed epigraphical splitting. In the following, Section 6.2 formulates the multiclass SVM learning problem with sparse regularization. Section 6.3 illustrates two algorithms for solving the considered problem with proximal tools. Finally, Section 6.4 evaluates the proposed approach on three standard datasets and compares it to state-of-the-art methods [20, 137, 140, 241].

6.2 Sparse Multiclass SVM

A multiclass classifier can be modeled as a function $d\colon \mathbb{R}^N \to \{1,\dots,K\}$ that predicts the class $k \in \{1,\dots,K\}$ associated to a given observation $u \in \mathbb{R}^N$ (e.g. a signal, an image or a graph). This predictor relies on K different *discriminating functions* $D_k\colon \mathbb{R}^N \mapsto \mathbb{R}$ which, for every $k \in \{1,\dots,K\}$, measure the likelihood that an observation belongs to the class k. Consequently, the predictor selects the class that best matches an observation, namely

$$d(u) = \operatorname*{arg\,max}_{k\in\{1,\dots,K\}} D_k(u). \tag{6.1}$$

The discriminating functions are built from a set of L input-output pairs

$$S = \big\{(u_\ell, z_\ell) \in \mathbb{R}^N \times \{1,\dots,K\} \mid \ell = \{1,\dots,L\}\big\}, \tag{6.2}$$

and they are assumed to be linear in some feature representation of inputs [76]. The latter assumption leads to the following form of the discriminating functions:

$$D_k(u) = \phi(u)^\top x^{(k)} + b^{(k)}, \tag{6.3}$$

where $\phi\colon \mathbb{R}^N \mapsto \mathbb{R}^M$ denotes a mapping from the input space onto an arbitrary feature space, and $(x^{(k)}, b^{(k)})_{1\le k\le K}$ denote the parameters to be estimated. For convenience, the latter ones are concatenated into a single vector $\mathrm{x} \in \mathbb{R}^{(M+1)K}$

$$\mathrm{x} = \begin{bmatrix} x^{(1)} \\ b^{(1)} \\ \vdots \\ x^{(K)} \\ b^{(K)} \end{bmatrix} \begin{matrix} \Big\} \mathrm{x}^{(1)} \\ \\ \\ \Big\} \mathrm{x}^{(K)} \end{matrix} \tag{6.4}$$

so that (6.3) can be shortened to $D_k(u) = \varphi(u)^\top \mathrm{x}^{(k)}$ with $\varphi(u) = \begin{bmatrix}\phi(u)^\top & 1\end{bmatrix}^\top$.

6.2.1 Background

The objective of learning is to find the vector x such that, for every $\ell \in \{1,\dots,L\}$, the input-output pair $(u_\ell, z_\ell) \in S$ is correctly predicted by the classifier, namely

$$z_\ell = \operatorname*{arg\,max}_{k\in\{1,\dots,K\}} \varphi(u_\ell)^\top \mathrm{x}^{(k)}. \tag{6.5}$$

By the definition of argmax, the above equality holds if

$$(\forall \ell \in \{1,\dots,L\}) \qquad \max_{k\neq z_\ell} \varphi(u_\ell)^\top (\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)}) < 0, \tag{6.6}$$

For simplicity, $k \neq z_\ell$ is short for $k \in \{1,\dots,K\} \setminus \{z_\ell\}$.

or, equivalently, with the introduction of a *margin* $\mu_\ell \in \,]0, +\infty[$,

$$(\forall \ell \in \{1,\dots,L\}) \qquad \max_{k\neq z_\ell} \varphi(u_\ell)^\top (\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)}) \le -\mu_\ell. \tag{6.7}$$

Unfortunately, this constraint has no practical interest for learning purposes, as it becomes infeasible when the training set is not fully separable. Multiclass SVMs overcome this issue by introducing the notion of *soft margins*, which consists of adding a vector of slack variables $\xi = (\xi^{(\ell)})_{1\leq\ell\leq L}$ into (6.7), yielding

$$\begin{cases} (\forall \ell \in \{1,...,L\}) & \max\limits_{k\neq z_\ell} \ \varphi(u_\ell)^\top(\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)}) \leq \xi^{(\ell)} - \mu_\ell, \\ (\forall \ell \in \{1,...,L\}) & \xi^{(\ell)} \geq 0, \end{cases} \tag{6.8}$$

The multiclass SVM learning problem is thus obtained by adding a quadratic regularization [78], leading to the convex optimization problem expressed as

The regularization does not involve the offsets $(b^{(k)})_{1\leq k\leq K}$.

$$\underset{(\mathrm{x},\xi)\in\mathbb{R}^{(M+1)K}\times\mathbb{R}^L}{\text{minimize}} \quad \sum_{k=1}^{K} \|x^{(k)}\|_2^2 + \lambda \sum_{\ell=1}^{L} \xi^{(\ell)} \qquad \text{s. t.}$$
$$\begin{cases} (\forall \ell \in \{1,...,L\}) & \max\limits_{k\neq z_\ell} \ \varphi(u_\ell)^\top(\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)}) \leq \xi^{(\ell)} - \mu_\ell, \\ (\forall \ell \in \{1,...,L\}) & \xi^{(\ell)} \geq 0, \end{cases} \tag{6.9}$$

where $\lambda \in]0, +\infty[$. Note that the linear penalty on the slack variables allows us to minimize the violation of constraint (6.7). By using standard convex analysis [24], the above problem can be equivalently rewritten without slack variables as

$$\underset{\mathrm{x}\in\mathbb{R}^{(M+1)K}}{\text{minimize}} \ \sum_{k=1}^{K} \|x^{(k)}\|_2^2 + \lambda \sum_{\ell=1}^{L} \max\left\{0, \mu_\ell + \max_{k\neq z_\ell} \ \varphi(u_\ell)^\top(\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)})\right\}. \tag{6.10}$$

Hereabove, the second term is called *hinge loss* when $\mu_\ell \equiv 1$.

6.2.2 Proposed approach

The proposed approach consists of replacing the squared $\boldsymbol{\ell}_2$-norm regularization in Problem (6.10) with a generic function $f \in \Gamma_0(\mathbb{R}^{(M+1)K})$. In addition, the hinge loss is rewritten in an equivalent form by introducing, for every $\ell \in \{1,\dots,L\}$, the linear operator $T_\ell \colon \mathbb{R}^{(M+1)K} \mapsto \mathbb{R}^K$ defined as

$$\left(\forall \mathrm{x} \in \mathbb{R}^{(M+1)K}\right) \qquad T_\ell\,\mathrm{x} = \begin{bmatrix} \varphi(u_\ell)^\top(\mathrm{x}^{(1)} - \mathrm{x}^{(z_\ell)}) \\ \vdots \\ \varphi(u_\ell)^\top(\mathrm{x}^{(K)} - \mathrm{x}^{(z_\ell)}) \end{bmatrix}, \tag{6.11}$$

the vector $r_\ell = (r_\ell^{(k)})_{1\leq k\leq K} \in \mathbb{R}^K$ defined as

$$(\forall k \in \{1,\dots,K\}) \qquad r_\ell^{(k)} = \begin{cases} 0, & \text{if } k = z_\ell, \\ \mu_\ell, & \text{otherwise,} \end{cases} \tag{6.12}$$

and the function $h_\ell \colon \mathbb{R}^K \mapsto \mathbb{R}$ defined, for every $\mathrm{y}^{(\ell)} = (y^{(\ell,k)})_{1\leq k\leq K} \in \mathbb{R}^K$, as

$$h_\ell(\mathrm{y}^{(\ell)}) = \max_{1\leq k\leq K} y^{(\ell,k)} + r_\ell^{(k)}, \tag{6.13}$$

so that the following holds

$$h_\ell(T_\ell \mathrm{x}) = \max\left\{0, \mu_\ell + \max_{k\neq z_\ell} \ \varphi(u_\ell)^\top(\mathrm{x}^{(k)} - \mathrm{x}^{(z_\ell)})\right\}. \tag{6.14}$$

The objective is to solve the following convex optimization problems

$$\textit{regularized approach:} \quad \underset{\mathrm{x}\in\mathbb{R}^{(M+1)K}}{\text{minimize}} \quad f(\mathrm{x}) + \lambda \sum_{\ell=1}^{L} h_\ell(T_\ell\,\mathrm{x}), \tag{6.15}$$

$$\textit{constrained approach:} \quad \underset{\mathrm{x}\in\mathbb{R}^{(M+1)K}}{\text{minimize}} \quad f(\mathrm{x}) \quad \text{s.t.} \quad \sum_{\ell=1}^{L} h_\ell(T_\ell\,\mathrm{x}) \leq \eta, \tag{6.16}$$

where λ and η are positive constants. Note that, by Lagrangian duality, the above formulations are equivalent for some specific values of η and λ. The interest of considering the constrained formulation lies in the fact that η may be easier to set, since it is directly related to the properties of the training data.

As mentioned in the introduction, the regularization term f is chosen so as to promote some form of sparsity. A popular example is the $\boldsymbol{\ell}_1$-norm, as it ensures that the solution will have a number of coefficients exactly equal to zero, depending on the strength of the regularization [9]. Another example is given by the mixed $\boldsymbol{\ell}_{1,p}$-norm. For every $\mathrm{x} \in \mathbb{R}^{(M+1)K}$, assume that, for each $k \in \{1,\ldots,K\}$, the vector $\mathrm{x}^{(k)} \in \mathbb{R}^{M+1}$ is block-decomposed as follows:

$$\mathrm{x}^{(k)} = \Big[\underbrace{\left(x^{(k,1)}\right)^\top}_{\text{size } M_1} \quad \ldots \quad \underbrace{\left(x^{(k,B)}\right)^\top}_{\text{size } M_B} \quad b^{(k)} \Big]^\top, \tag{6.17}$$

with $M_1 + \cdots + M_B = M$. The $\boldsymbol{\ell}_{1,p}$-norm is defined as

$$f(\mathrm{x}) = \sum_{k=1}^{K} \sum_{b=1}^{B} \|x^{(k,b)}\|_p. \tag{6.18}$$

The mixed-norm regularization is known to induce *block-sparsity*: the solution is partitioned into groups and the components of each group are ideally either all zeros or all non-zeros. In this context, the exponent values $p = 2$ or $p = +\infty$ are the most popular choices. In particular, the $\boldsymbol{\ell}_{1,\infty}$-norm tends to favor solutions with few nonzero groups having components of similar magnitude.

6.2.3 Connection with logistic regression

Problems (6.15)-(6.16) can model different types of learning if the hinge loss defined in (6.13) is replaced with other functions, such as the quadratic loss, the exponential loss, the squared hinge loss, the modified Huber loss, or the logistic loss (e.g., see [242]). These functions actually provide a *smooth approximation* of the hinge loss. On the other side, the hinge loss can be interpreted as a piecewise-affine lower approximation of the logistic loss, as [24]

$$\max_{1\leq k\leq K} y^{(\ell,k)} \leq \log\Big(\sum_{k=1}^{K} \exp\left(y^{(\ell,k)}\right)\Big) \leq \max_{1\leq k\leq K} y^{(\ell,k)} + \log(K), \tag{6.19}$$

where the second inequality is tight in the case when $y^{(\ell,1)} = \cdots = y^{(\ell,K)}$. Consequently, the hinge-loss constraint in (6.16) can be regarded as an outer approximation of the logistic-loss constraint in Problem (4.44). In this respect, the algorithm for solving Problem (6.16) proposed in Section 6.3.2 can be seen as an efficient implementation of the approximation technique derived in Chapter 4.

6.3 Optimization method

Two different algorithms are now proposed for solving Problems (6.15) and (6.16).

6.3.1 Regularized formulation

Problem (6.15) fits nicely into the framework provided by FBPD method [69], as the proximity operator of both f and $(h_\ell)_{1\leq\ell\leq L}$ can be efficiently computed. Indeed, prox_f has a closed form for several norms [65, 71, 178], while $(\mathrm{prox}_{h_\ell})_{1\leq\ell\leq L}$ can be computed by the projection onto the standard simplex, as described next.

Proposition 6.3.1. *For every $\ell \in \{1,\dots,L\}$,*

$$(\forall y^{(\ell)} \in \mathbb{R}^K) \quad \mathrm{prox}_{\lambda h_\ell}\big(y^{(\ell)}\big) = y^{(\ell)} - P_{S_\lambda}\big(y^{(\ell)} + r_\ell\big), \tag{6.20}$$

with

$$S_\lambda = \Big\{u = \big(u^{(k)}\big)_{1\leq k\leq K} \in [0,+\infty[^K \;\Big|\; \sum_{k=1}^K u^{(k)} = \lambda\Big\}. \tag{6.21}$$

Proof. Note that $u \in \mathbb{R}^K \mapsto \lambda \max_{1\leq k\leq K} u^{(k)}$ is the support function of S_λ, defined as $(\forall u \in \mathbb{R}^K)\ \sigma_{S_\lambda}(u) = \sup_{v\in S_\lambda} v^\top u$. Hence, for every $y^{(\ell)} \in \mathbb{R}^K$, $\lambda h_\ell(y^{(\ell)}) = \sigma_{S_\lambda}(y^{(\ell)} + r_\ell)$ and

$$\mathrm{prox}_{\lambda h_\ell}\big(y^{(\ell)}\big) = \mathrm{prox}_{\sigma_{S_\lambda}}(y^{(\ell)} + r_\ell) - r_\ell, \tag{6.22}$$

Since σ_{S_λ} is the conjugate function of ι_{S_λ}, (6.20) is deduced by applying Moreau's decomposition formula [12, Theorem 14.3(ii)]. □

If the function h_ℓ is replaced with the logistic loss in (4.45)*, the term $-r_\ell^\top y^{(\ell)}$ in the dual formulation is replaced by*

$$\sum_{k=1}^K y^{(\ell,k)} \log(y^{(\ell,k)})$$

with the convention $0\log 0 = 0$.

Problem (6.15) are summarized in Algorithm 6.1, with the notation $T = [T_1^\top \ \dots \ T_L^\top]^\top$ and $r = [r_1^\top \ \dots \ r_L^\top]^\top$. As mentioned in Section 1.2.4, this algorithm solves both Problem (6.15) and its dual formulation

$$\underset{y\in\mathbb{R}^{LK}}{\text{minimize}} \quad f^*(-T^\top y) - \sum_{\ell=1}^L r_\ell^\top y^{(\ell)} \quad \text{s.t.} \quad (\forall \ell \in \{1,...,L\}) \quad y^{(\ell)} \in S_\lambda. \tag{6.23}$$

In the case when $f = (1/2)\|\cdot\|_2^2$, the primal and dual solutions are linked by $\mathrm{x} = -T^\top y$, and thus Problem (6.23) reduces to the Lagrangian dual formulation of Problem (6.9) used in standard SVMs [19, 78]. The projection onto the simplex can be efficiently computed with the method illustrated in Section 1.3.1.

Algorithm 6.1 *FBPD for solving Problem* (6.15)

Initialization

choose $(\mathrm{x}^{[0]}, y^{[0]}) \in \mathbb{R}^{(M+1)K} \times \mathbb{R}^{LK}$

set $\tau > 0$ and $\sigma > 0$ such that $\tau\sigma\|T\|^2 \leq 1$

For $i = 0, 1, \dots$

$\mathrm{x}^{[i+1]} = \mathrm{prox}_{\tau f}\left(\mathrm{x}^{[i]} - \tau\, T^\top y^{[i]}\right)$

$y^{[i+1]} = P_{(S_\lambda)^L}\Big(y^{[i]} + \sigma T\big(2\mathrm{x}^{[i+1]} - \mathrm{x}^{[i]}\big) + \sigma\, r\Big)$

6.3.2 Constrained formulation

The resolution of Problem (6.16) presents a challenging computational issue, as the projection onto the hinge-loss constraint set cannot be computed in closed form. One way to manage this constraint consists of introducing an auxiliary vector $\zeta = \big(\zeta^{(\ell)}\big)_{1\leq\ell\leq L}$, so that Problem (6.16) can be equivalently rewritten as

$$\underset{(\mathrm{x},\zeta)\in\mathbb{R}^{(M+1)K}\times\mathbb{R}^L}{\text{minimize}} \quad f(\mathrm{x}) \quad \text{s. t.} \quad \begin{cases} \displaystyle\sum_{\ell=1}^{L}\zeta^{(\ell)} \leq \eta, \\ (\forall \ell \in \{1,\ldots,L\}) \quad h_\ell(T_\ell\,\mathrm{x}) \leq \zeta^{(\ell)}. \end{cases} \tag{6.24}$$

While this approach is conceptually similar to adding the slack variables in (6.8), the epigraphical splitting specifically aims at simplifying the way of solving the problem. A possible interpretation of Problem (6.24) is indeed the following

$$\underset{(\mathrm{x},\zeta)\in\mathbb{R}^{(M+1)K}\times\mathbb{R}^L}{\text{minimize}} \quad f(\mathrm{x}) \qquad \text{s. t.} \qquad \begin{cases} (T\mathrm{x},\,\zeta) \in E, \\ \zeta \in V, \end{cases} \tag{6.25}$$

where

$$E = \big\{(y,\zeta)\in\mathbb{R}^{LK}\times\mathbb{R}^L \;\big|\; (\forall\ell\in\{1,\ldots,L\}) \quad (y^{(\ell)},\zeta^{(\ell)})\in \operatorname{epi} h_\ell\big\}, \tag{6.26}$$

$$V = \big\{\zeta\in\mathbb{R}^L \;\big|\; 1_L^\top\zeta\leq\eta\big\}. \tag{6.27}$$

The iterations of FBPD related to Problem (6.25) are listed in Algorithm 6.2. The advantage of this approach is that the projection P_E is given by Proposition 2.3.8.

Algorithm 6.2 *FBPD for solving Problem* (6.16)

Initialization

$$\begin{bmatrix} \text{choose } (\mathrm{x}^{[0]},\zeta^{[0]})\in\mathbb{R}^{(M+1)K}\times\mathbb{R}^L \\ \text{choose } (y^{[0]},\xi^{[0]})\in\mathbb{R}^{L(K-1)}\times\mathbb{R}^L \\ \text{set } \tau>0 \text{ and } \sigma>0 \text{ such that } \tau\sigma\max\{\|T\|^2,1\}\leq 1. \end{bmatrix}$$

For $i = 0, 1, \ldots$

$$\begin{bmatrix} \mathrm{x}^{[i+1]} = \operatorname{prox}_{\tau f}\big(\mathrm{x}^{[i]} - \tau\, T^\top y^{[i]}\big) \\ \zeta^{[i+1]} = P_V\big(\zeta^{[i]} - \tau\,\xi^{[i]}\big) \\ \widehat{y}^{[i]} = y^{[i]} + \sigma T\big(2\mathrm{x}^{[i+1]} - \mathrm{x}^{[i]}\big) \\ \widehat{\xi}^{[i]} = \xi^{[i]} + \sigma\big(2\zeta^{[i+1]} - \zeta^{[i]}\big) \\ \big(\widetilde{y}^{[i]},\widetilde{\xi}^{[i]}\big) = P_E\big(\widehat{y}^{[i]}/\sigma,\widehat{\xi}^{[i]}/\sigma\big) \\ y^{[i+1]} = \widehat{y}^{[i]} - \sigma\widetilde{y}^{[i]} \\ \xi^{[i+1]} = \widehat{\xi}^{[i]} - \sigma\widetilde{\xi}^{[i]}. \end{bmatrix}$$

Table 6.1 *Comparisons on the leukemia database.*

$f(\mathrm{x})$	HINGE		SQUARE		LOGIT		ONE-VS-ALL	
	errors	non-zero coeff.	errors	non-zero coeff.	errors	non-zero coeff.	errors	non-zero coeff.
$\boldsymbol{\ell}_2$	**1/34**	7129 + 7129 + 7129	2/34	7129 + 7129 + 7129	**1/34**	7129 + 7129 + 7129	2/34	7129 + 7129 + 7129
$\boldsymbol{\ell}_1$	**2/34**	13 + 03 + 10	3/34	8 + 3 + 8	3/34	18 + 05 + 14	3/34	19 + 8 + 15
$\boldsymbol{\ell}_{1,2}$	**0/34**	95 + 5 + 75	1/34	55 + 05 + 45	**0/34**	50 + 05 + 35	1/34	70 + 10 + 50
$\boldsymbol{\ell}_{1,\infty}$	**0/34**	50 + 5 + 45	**0/34**	35 + 05 + 35	**0/34**	50 + 05 + 40	**0/34**	45 + 5 + 45

6.4 Numerical results

The performance of sparse multiclass SVM is evaluated on the following databases.

- **Leukemia database.** The first experiment concerns the classification of microarray data. The considered database contains 72 samples of $N = M = 7129$ gene expression levels (so that $\phi(u) = u$) measured from patients having $K = 3$ types of leukemia disease [112]. The database is usually organized in $L = 38$ training samples and 34 test samples.[1] In the experiments, the mixed-norm regularization is defined on blocks of size 5.

- **MNIST dataset.** The second experiment concerns the classification of handwritten digits [142]. The MNIST database contains a number of 28×28 grayscale images ($N = 784$) displaying digits from 0 to 9 ($K = 10$). The database is organized in 60000 training images and 10000 test images.[2] In the experiments, the mapping ϕ is defined through the scattering convolution network [31] with $\overline{m} = 2$ wavelet layers scaled up to $2^J = 4$, which transforms an input image of size 28×28 in 81 images of size 14×14 (thus $M = 15876$). For this database, indeed, it was observed [31] that the scattering network leads to better results than kernel PCA [197] or deep convolution networks [143, 144]. For the regularization, the $\boldsymbol{\ell}_{1,\infty}$-norm is defined by dividing each vector $(\mathrm{x}^{(k)})_{1 \leq k \leq K}$ in 14^2 blocks of size 81. Moreover, in order to evaluate the performance, the classifier is trained on 25 different training subsets of size $L \in \{3K, 5K, 10K\}$, the classification errors are computed by evaluating the 25 trained classifiers on the whole test set, and by averaging the resulting errors.

- **News20 database.** The third experiment concerns the classification of text documents into a fixed number of predefined categories [139]. The News20 database contains a number of documents partitioned across $K = 20$ different newsgroups. The database is organized in 11314 training documents and 7532 test documents.[3] In the experiments, the mapping ϕ is defined through the *term frequency – inverse document frequency* transformation [125], yielding $M = 26214$. For the regularization, the $\boldsymbol{\ell}_{1,2}$-norm is defined so as to regroup the same feature across all classes, in the same way as [20]. Moreover, in order to evaluate the performance, the classifier is trained on 10 different training subsets of size $L \in \{5K, 10K, 50K\}$, the classification errors are computed by evaluating the 10 trained classifiers on the whole test set, and by averaging the resulting errors.

[1] Data available at *www.broadinstitute.org/cancer/software/genepattern/datasets*
[2] Data available at *http://yann.lecun.com/exdb/mnist*
[3] Data available at *www.cad.zju.edu.cn/home/dengcai/Data/TextData.html*

6.4.1 Assessment of classification accuracy

This section evaluates the classification errors obtained with the sparse multiclass SVM formulated in Problems (6.15)-(6.16). The objective here is to show that the *exact* hinge loss allows one to achieve better performance than its approximated smooth versions, especially with a few training data. Hence, the proposed method is compared with the following approaches.

- The multiclass SVM proposed by Blondel et al. [20]

$$\underset{\mathrm{x}\in\mathbb{R}^{(M+1)K}}{\text{minimize}} \;\; f(\mathrm{x})+\lambda\sum_{\ell=1}^{L}\sum_{k\neq z_\ell}\Big(\max\left\{0,\mu_\ell+\varphi(u_\ell)^\top(\mathrm{x}^{(k)}-\mathrm{x}^{(z_\ell)})\right\}\Big)^2. \quad (6.28)$$

- The multinomial logistic regression (e.g., see [137])

$$\underset{\mathrm{x}\in\mathbb{R}^{(M+1)K}}{\text{minimize}} \;\; f(\mathrm{x})+\lambda\sum_{\ell=1}^{L}\log\Big(1+\sum_{k\neq z_\ell}\exp\Big\{\mu_\ell+\varphi(u_\ell)^\top(\mathrm{x}^{(k)}-\mathrm{x}^{(z_\ell)})\Big\}\Big). \quad (6.29)$$

- The binary SVM by Laporte et al. [140] based on the "one-vs-all" strategy. For every $k\in\{1,\ldots,K\}$, the latter amounts to separately estimating the vector $\mathrm{x}^{(k)}$ by splitting the training samples $(u_\ell,\widetilde{z}_\ell)$ in two classes, with $\widetilde{z}_\ell=1$ if $z_\ell=k$, and $\widetilde{z}_\ell=-1$ otherwise, thus leading to

$$\underset{\mathrm{x}^{(k)}\in\mathbb{R}^{(M+1)}}{\text{minimize}} \;\; f(\mathrm{x})+\lambda\sum_{\ell=1}^{L}\Big(\max\left\{0,\mu_\ell+\widetilde{z}_\ell\,\varphi(u_\ell)^\top\mathrm{x}^{(k)}\right\}\Big)^2, \quad (6.30)$$

Note that (6.30) *is the specialization of* (6.28) *to the case* $K=2$*, after a variable change.*

Problems (6.15)-(6.16) are referred to as *hinge*, whereas Problems (6.28)-(6.30) are referred to as *square*, *logit*, and *one-vs-all*. Since the parameters λ and η need to be estimated (e.g., through cross validation), it is important to evaluate the impact of their choice on the performance. To compare the above methods for different choices of these parameters, it was set $\lambda=\alpha^{-1}$ or $\eta=\alpha L$, with α being varied inside a fixed set of predefined values. It was also set $\mu_\ell\equiv 1$.

- **Leukemia database**. Table 6.1 reports the classification errors and the number of non-zero coefficients in $(\mathrm{x}^{(k)})_{1\leq k\leq 3}$ obtained with *hinge*, *square*, *logit*, and *one-vs-all* using various regularization terms. For each method, we set α so as to achieve the best accuracy (by using a simple trial-and-error strategy). The results indicate that sparse regularization effectively selects a small set of important features for each prediction vector $(\mathrm{x}^{(k)})_{1\leq k\leq 3}$, with better results than the quadratic regularization. In addition, the classification errors show that *hinge* is often more accurate than *square*.

- **MNIST database**. Figures 6.1a, 6.1c and 6.1e report the classification errors as a function of the regularization hyperparameter. These results were obtained with the $\boldsymbol{\ell}_{1,\infty}$-norm regularization, as it was the one leading to the best results in all our experiments on this database. The classification errors indicate that the *hinge* approach is slightly more accurate than the other ones. On the other side, Figures 6.1b, 6.1d and 6.1f report the percentage of zero coefficients in vectors $(\mathrm{x}^{(k)})_{1\leq k\leq K}$ as a function of α. The plots show that the *hinge* approach yields solutions slightly more sparse than the other ones.

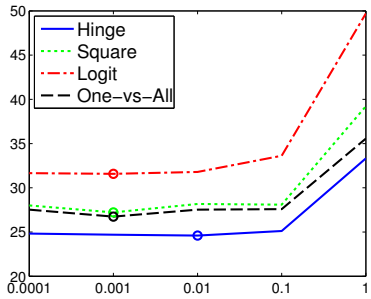

(a) $L/K = 3$ (errors vs α)

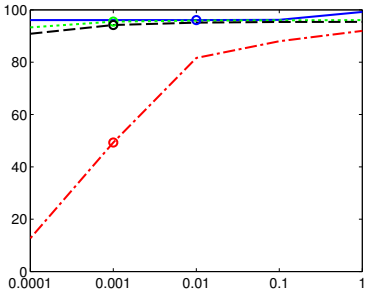

(b) $L/K = 3$ (sparsity vs α)

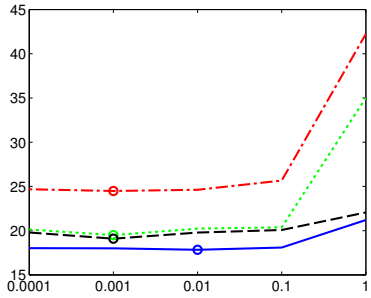

(c) $L/K = 5$ (errors vs α)

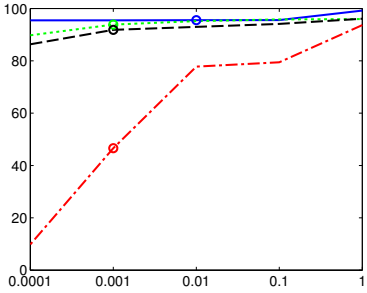

(d) $L/K = 5$ (sparsity vs α)

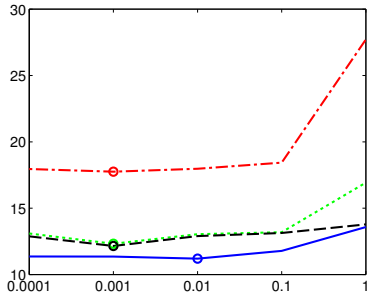

(e) $L/K = 10$ (errors vs α)

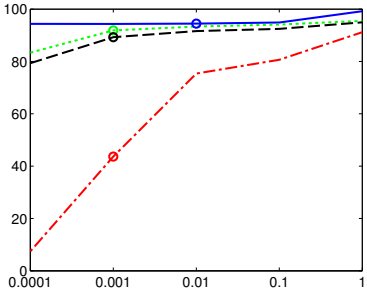

(f) $L/K = 10$ (sparsity vs α)

Figure 6.1 *Results on MNIST database with the $\ell_{1,\infty}$-regularization for $L \in \{3K, 5K, 10K\}$. Left column: classification errors as a function of α. Right column: percentage of zero coefficients in vectors $(\mathrm{x}^{(k)})_{1\leq k\leq K}$ as a function of α. The circles mark the values yielding the best accuracy.*

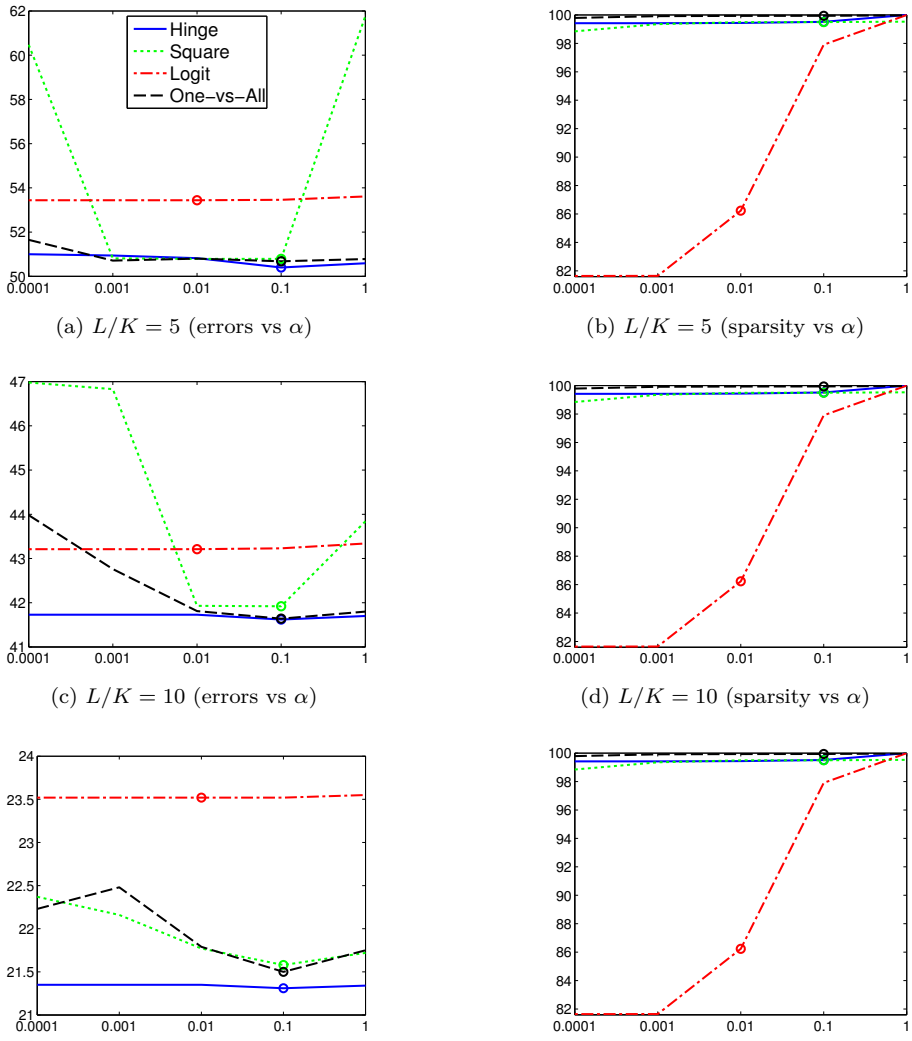

(a) $L/K = 5$ (errors vs α)

(b) $L/K = 5$ (sparsity vs α)

(c) $L/K = 10$ (errors vs α)

(d) $L/K = 10$ (sparsity vs α)

(e) $L/K = 50$ (errors vs α)

(f) $L/K = 50$ (sparsity vs α)

Figure 6.2 *Results on News20 database with the $\ell_{1,2}$-regularization for $L \in \{5K, 10K, 50K\}$. Left column: classification errors as a function of α. Right column: percentage of zero coefficients in vectors $(\mathrm{x}^{(k)})_{1 \leq k \leq K}$ as a function of α. The circles mark the values yielding the best accuracy.*

- **News20 database**. Figures 6.2a, 6.2c and 6.2e report the classification errors (as a function of the regularization hyperparameter) obtained by using the $\boldsymbol{\ell}_{1,2}$-norm regularization. The classification errors indicate that the *hinge* approach is slightly more accurate than the *square* approach. The plots also show that the results obtained with the hinge approach are more robust w.r.t. the choice of the regularization parameter. On the other side, Figures 6.2b, 6.2d and 6.2f report the percentage of zero coefficients in vectors $(\mathrm{x}^{(k)})_{1\leq k\leq K}$ as a function of α. The plots show that the *hinge* approach yields solutions as sparse as the *square* approach.

6.4.2 Assessment of execution times

In this section, the execution time of Algorithms 6.1 and 6.2 is compared with

- a FISTA implementation of Problem (6.28),
- a forward-backward implementation of Problem (6.29),
- a FBPD implementation of Problem (6.24) being reformulated with an extended number of linear constraints

$$\underset{(\mathrm{x},\zeta)\in\mathbb{R}^{(M+1)K}\times\mathbb{R}^{LK}}{\text{minimize}} \quad f(\mathrm{x}) \quad \text{s. t.}$$
$$\left\{\begin{array}{ll} \displaystyle\sum_{\ell=1}^{L}\sum_{k=1}^{K}\zeta^{(\ell,k)} \leq K\eta, & \\ (\forall \ell\in\{1,...,L\}) & \zeta^{(\ell,1)} = \cdots = \zeta^{(\ell,K)}, \\ (\forall \ell\in\{1,...,L\}) & \zeta^{(\ell,1)} \geq 0, \ldots, \zeta^{(\ell,K)} \geq 0, \\ (\forall \ell\in\{1,...,L\}) & T_\ell\,\mathrm{x} + r_\ell - (\zeta^{(\ell,k)})_{1\leq k\leq K} \leq 0. \end{array}\right. \tag{6.31}$$

  This approach is conceptually similar to the linear programming methods proposed by [220] and [241] for $\boldsymbol{\ell}_1$- or $\boldsymbol{\ell}_{1,+\infty}$-regularized SVMs.

Figures 6.3a, 6.3c and 6.3e show the execution times (averaged among 10 training sets) obtained by the above algorithms for various values of λ and η on the MNIST database with $L\in\{3K, 5K, 10K\}$. In this experiment, the execution times refer to a stopping criterion of 10^{-5} on the relative error between two consecutive iterates. Figures 6.3b, 6.3d and 6.3f show the relative distance to $\|\mathrm{x}^{[i]} - \mathrm{x}^{[\infty]}\|/\|\mathrm{x}^{[\infty]}\|$ (as a function of time) for the values of λ and η yielding the best accuracy (as reported in Figure 6.1), where $\mathrm{x}^{[\infty]}$ denotes the solution computed with a stopping criterion of 10^{-5}. These results demonstrate that the proposed algorithms are faster than the approaches based on linear constraints and logistic regression, while being comparable in terms of execution times to approaches based on the square hinge loss. In addition, Algorithm 6.2 turns out to converge faster than Algorithm 6.1. This can be explained by the higher computational cost of the projection onto the standard simplex. The codes were tested in Matlab on a Intel CPU at 3.33 GHz and 24 GB of RAM.

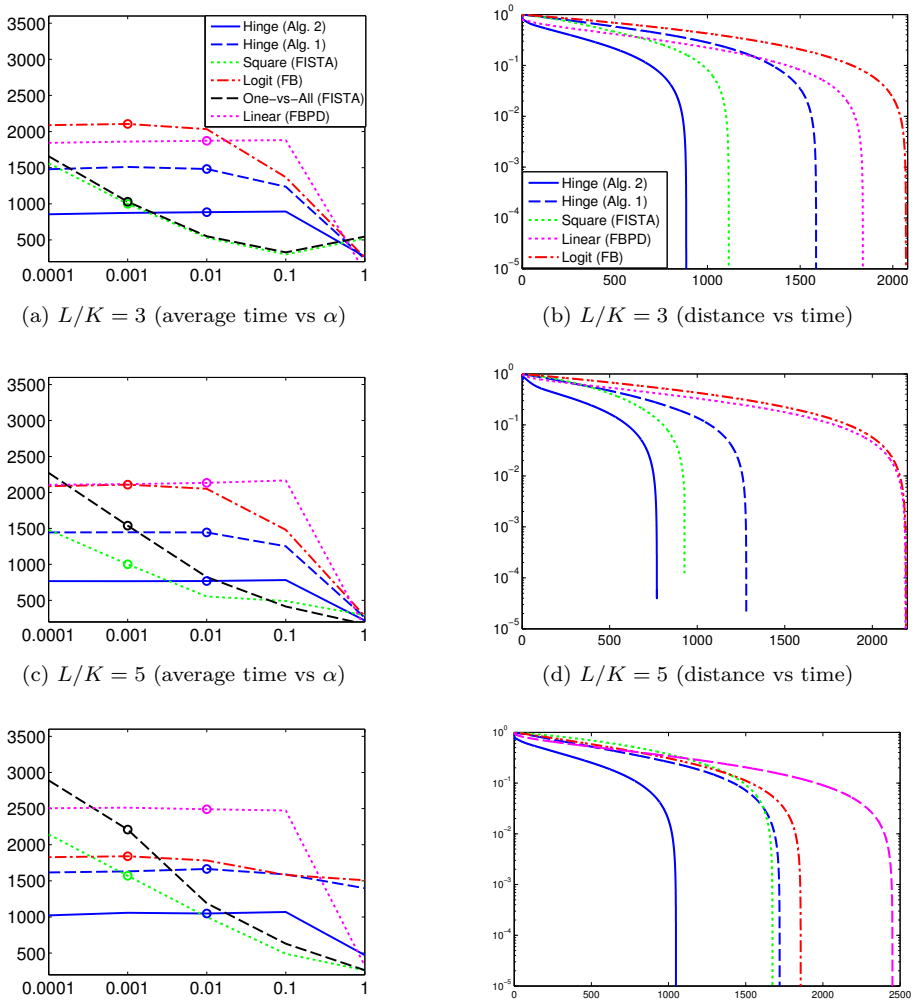

(a) $L/K = 3$ (average time vs α)

(b) $L/K = 3$ (distance vs time)

(c) $L/K = 5$ (average time vs α)

(d) $L/K = 5$ (distance vs time)

(e) $L/K = 10$ (average time vs α)

(f) $L/K = 10$ (distance vs time)

Figure 6.3 *Results on MNIST database with the $\ell_{1,\infty}$-regularization for $L \in \{3K, 5K, 10K\}$. Left column: execution time as a function of α, where the circles mark the values yielding the best accuracy (as reported in Figure 6.1). Right column: distance to $\mathrm{x}^{[\infty]}$ (as a function of time) obtained with the values of α marked by a circle in the left column (note that the one-vs-all approach, being defined by multiple optimization problems, does not allow us to determine the iterate $x^{[i]}$ at each iteration, hence the associated plot cannot be traced).*

6.4.3 Quadratic regularization

Although the emphasis is on sparse learning, we propose to complete our analysis by evaluating the efficiency of the proposed algorithms in the case when f is a quadratic regularization function. To this end, the execution times of Algorithms 6.1 and 6.2 is compared with the *SVM-struct* algorithm proposed by Joachims et al. [126], which provides a numerical approach for solving Problem (6.9) through a cutting-plane technique. Figure 6.4 reports the execution times (averaged on 10 training sets) obtained by the above methods on the MNIST database with $L \in \{3K, 5K, 10K, 50K, 100K, 500K\}$ and different values of α. In this experiment, the stopping criterion was set to 10^{-3} in all methods, and the regularization parameter of *SVM-struct* was set to L/α.

The results show that the proposed algorithms are competitive with state-of-the-art solutions in scenarios with a limited number of training data. The same cannot be claimed for larger databases, as *SVM-struct* scales particularly well w.r.t. the number M of features and the size L of the training set. Note however that, when $L/K = 500$, the number of significant features for the SVM classifier designed with a quadratic regularization is equal to 158214 (w.r.t. a threshold set to 10^{-5}), while a sparse approach using an $\ell_{1,\infty}$-norm regularization yields only 42795 nonzero features.

6.5 Conclusions

We have proposed two efficient algorithms for learning a sparse multiclass SVM. Our approach makes it possible to minimize a criterion involving the multiclass hinge loss and a sparsity-inducing regularization. In the literature, such a criterion is typically approximated by replacing the hinge loss with a smooth penalty, such as the quadratic hinge loss or the logistic loss. In this chapter, we have provided two solutions that directly deal with the hinge loss: one addressing the regularized formulation and the other one adapted to the constrained formulation. The performance of the proposed solutions have been evaluated over three databases in scenarios with a few training data. The results show that the use of the hinge loss, rather than an approximation, leads to a slightly better classification accuracy and tends to make the method more robust w.r.t. the choice of the regularization parameter, while the proposed algorithms are often faster than state-of-the-art solutions.

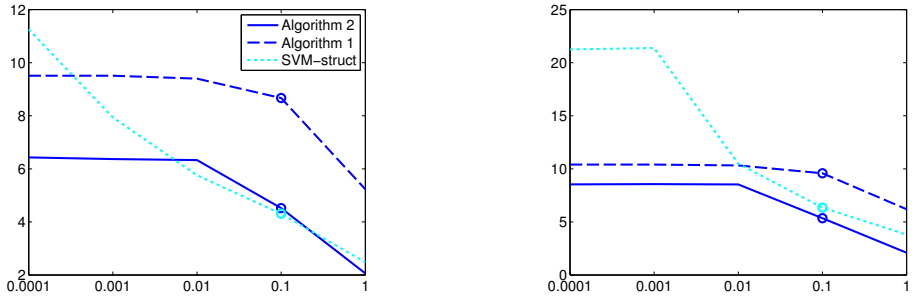

(a) $L/K = 3$ (average time vs α)

(b) $L/K = 5$ (average time vs α)

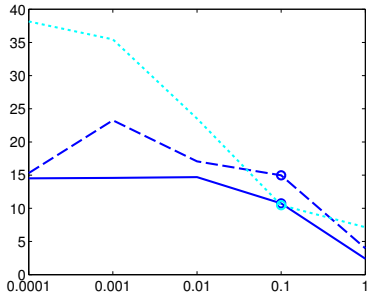

(c) $L/K = 10$ (average time vs α)

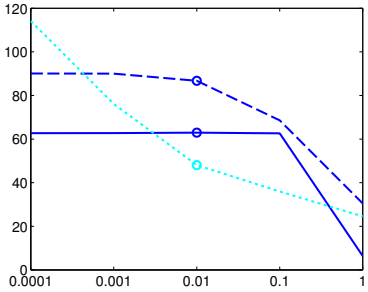

(d) $L/K = 50$ (average time vs α)

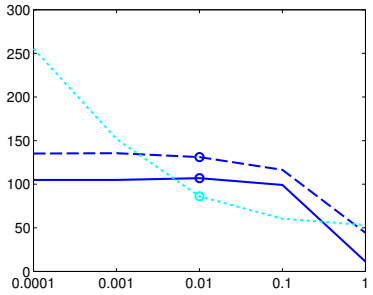

(e) $L/K = 100$ (average time vs α)

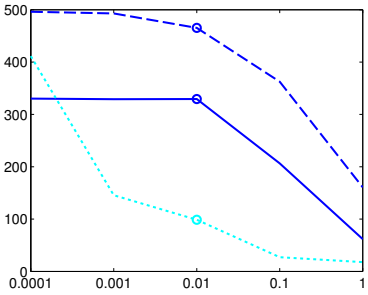

(f) $L/K = 500$ (average time vs α)

Figure 6.4 *Results on MNIST database with a quadratic regularization. The plots show the execution times obtained with a stopping criterion of 10^{-3} for some values of α. The circles mark the values of α yielding the best classification accuracy.*

The saddest aspect of life right now is that science gathers knowledge faster than society gathers wisdom.

Isaac Asimov

Chapter 7
Image Forgery Detection

The main focus of this chapter is the detection of image forgeries with the use of photo response non-uniformity, a deterministic pattern noise specific of each individual camera. To this end, we propose a variational approach that aims at solving a binary segmentation problem regularized with a ℓ_1-norm constraint, and we explain how to tackle it by convex relaxation and the proposed epigraphical splitting. Experiments on real and simulated forgeries show the good performance of our approach.

7.1 Introduction

Digital images are more and more frequently used to support important decisions. This is especially true in the forensic field where, to make just a few examples, images are routinely used to describe the scene of a crime, or to define responsibilities in road accidents. Unfortunately, with the wide availability of sophisticated image manipulation tools, modifying a digital photo with little or no obvious signs of tampering has become easier than ever before (e.g., see `www.fourandsix.com/photo-tampering-history`). Therefore, it is important to devise tools that help deciding on the authenticity of a digital image, which raises attention on the image forgery detection field.

Several approaches have been proposed in the literature to detect image alterations under a variety of scenarios. A first category comprises active techniques for image authentication based on the use of watermarks [244] and signatures [10, 243], with the former being embedded into images (possibly originating small distortions), and the latter being attached to images as side information. Although these methods are very effective, they can be applied only when the digital source is protected at the origin, which is probably a minority of the cases of interest. Therefore, there has been a steadily growing interest on passive techniques, which retrieve traces of manipulations from the image itself with no need of collaboration on the part of the user.

Some techniques are specifically tailored to copy-move forgeries, where portions of the image are cut and pasted elsewhere in the same image. Duplicated parts are discovered by block-based processing or, more efficiently, by means of suitable invariant features [5, 56, 130]. A more general approach considers physical inconsistencies, such as the lighting of objects, shadows, or geometric features (dimension, position, etc.) of objects w.r.t. the camera [127, 128, 154]. Also, as many images are saved in some compressed format, several forgery detection techniques rely on the traces left by multiple compressions. In fact, when a JPEG image is modified and saved again in JPEG format, specific artifacts appear as a result of the multiple quantization processes, suggesting the presence of some form of tampering [17, 44, 152, 237].

Another valuable source of information is the acquisition phase, which often

leaves peculiar traces, related to lens characteristics [101, 231], the color filter array (CFA) pattern [97, 136, 185], or the sensor array [41, 160], that can be used to discover image manipulations. In the latter context, the photo-response non uniformity (PRNU) noise appears as one of the most promising tools at hand. The PRNU arises from tiny imperfections in the silicon wafer used to manufacture the imaging sensor [114]. These physical differences provide a unique sensor pattern, specific of each individual camera, constant in time, and independent of the scene. It can be therefore considered as a sort of camera fingerprint and used as such to accomplish forgery detection or image identification tasks. Indeed, the most common forms of image forgery, like copy-move or splicing, delete the original camera PRNU from the target region, a fact that can be detected through suitable analyses, provided the camera PRNU is available. Note that, unlike with most other approaches, the detection of tampering is based on the *absence* of the fingerprint, hence does not depend on the specific type of forgery. On the other hand, the PRNU pattern is fairly robust to several forms of image processing, such as JPEG compression, filtering, or gamma correction [41, 160].

The main contribution of this chapter is the design of a new algorithm for PRNU-based forgery detection. The proposed strategy improves upon the reference method [41] in several aspects. The most important one consists of formulating the problem in terms of a binary segmentation regularized with a suitable Markov random field. So doing, one can exploit the strong spatial dependencies of the source, and take the decisions jointly on the whole image, rather than individually for each pixel.

7.1.1 Related work

A PRNU-based technique for camera identification and forgery detection was originally proposed in [160], then refined in [41]. Given the potential of this algorithm [43], many research groups have soon started working, trying to improve the estimate of the image PRNU, to define better decision statistics, or better decision strategies. In fact, since the PRNU is a very weak signal, its reliable estimation is crucial for the algorithm success.

An estimate of the PRNU is typically computed by subtracting a filtered version of the image from the observed one, obtaining a residual where the PRNU is present, but the image (seen as noise in this context) is mostly removed. However, the residual contains also traces of the signal, especially at high frequencies, due to the imperfection of the filtering process or to in-camera processing, such as JPEG compression, CFA interpolation, or vignetting [146].

In order to mitigate the aforementioned issue, the denoising filter used in the original technique was replaced with a state-of-the-art nonlocal filter [47], with significant performance improvements. A different strategy to reduce the interference of scene was proposed in [147], based on selective attenuation of wavelet transform coefficients.

Another major challenge is the suppression of the so-called non-unique artifacts [100], specific to a camera model or manufacturer. These include for example JPEG block artifacts, and CFA interpolation artifacts, both characterized by regular "linear" spatially periodic patterns, relatively easy to correct [148]. Non-unique artifacts may lead to wrong results, especially in camera identification, because of the increased similarity between the PRNU fingerprints

of a different devices with similar characteristics. Recently, nonlinear artifacts have also been reported, due to correction of radial lens distortion [110] and other advanced in-camera processing procedures [106].

As for decision statistics, since the normalized correlation used in the original algorithm was very sensitive to artifacts, the more stable peak-to-correlation energy (PCE) ratio was proposed in [100, 109] for camera identification purposes, further modified in [131] to lower the false positive rate. A different approach proposed in [239] was the adoption of canonical correlation analysis. In [48], it was proposed to compute statistics, and possibly take decisions, based on a prior segmentation of the image, thus moving towards an object-oriented processing. The original signal was used also in [153], where only regions characterized by higher signal quality are used, discarding those regions heavily deteriorated by irrelevant noise. Some papers, finally, focus on computational/storage complexity, certainly an issue for applications that involve large-scale databases, proposing the use of a quantized [13], or spatially limited [111], or hashed [120] PRNU.

7.1.2 Contributions

This chapter proposes a new PRNU-based forgery detection algorithm that relies on the same general structure and basic tools as the reference method [41], but improves upon it under several respects:

(i). first of all, the constant false alarm rate decision strategy is abandoned in favor of a more flexible Bayesian rule;

(ii). more important, decisions are now made jointly on the whole image, rather than individually for each pixel;

(iii). to do so, the strong spatial dependencies of the source are taken into account by modeling the data through a suitable Markov random field;

(iv). this leads to a binary segmentation problem involving a regularity constraint based on the $\boldsymbol{\ell}_1$-norm, which is tackled by convex relaxation and solved through the proposed epigraphical splitting, so as to guarantee the convergence to a global optimum in a limited time;

(v). moreover, the quality of the observed data is improved by using a nonlocal denoising algorithm;

(vi). finally, experiments prove that the proposed algorithm outperforms the reference one, with a very limited increase in the computational burden.

7.1.3 Outline

The chapter is organized as follows. Section 7.2 thoroughly revises the reference method [41]. Section 7.3 motivates and describes in detail the proposed improvements. Section 7.4 analyzes the performance by means of simulation experiments. Finally, Section 7.5 draws conclusions and outlines future research.

7.2 Background

Let $y \in \mathbb{R}^N$ be a digital image observed at the camera output (either as a single color band or the composition of multiple color bands), with $y^{(\ell)}$ denoting the pixel value at site ℓ. According to the standard camera model [41, 114], an image y can be written as

$$(\forall \ell \in \{1, \dots, N\}) \qquad y^{(\ell)} = \Big(1 + k^{(\ell)}\Big) x^{(\ell)} + \theta^{(\ell)}, \tag{7.1}$$

where $x \in \mathbb{R}^N$ is the ideal noise-free image, $k \in \mathbb{R}^N$ the camera PRNU, and $\theta \in \mathbb{R}^N$ the realization of an additive noise term which accounts for all types of disturbances. Let $\circ$ denotes the element-wise product. By rewriting (7.1) as

$$y = x \circ k + x + \theta, \tag{7.2}$$

we stress from the beginning that the PRNU k is the only signal of interest in all our analyses, and all other terms assume the role of unwanted disturbance. This includes the ideal image x, which is in fact estimated and subtracted from the original image, obtaining a more tractable *noise residual*

$$r = y - \widehat{x}, \tag{7.3}$$

where $\widehat{x} = f(y)$ is estimated by means of a denoising filter f based on a wavelet transformation [170]. Even so, since the PRNU is typically very small, except for possible faulty sensors, we work in a very hostile environment, with an extremely low signal-to-noise (SNR) ratio.

For convenience, the above residual is rewritten so that k multiplies the observed image y rather than the unknown original x, leading to

$$\begin{aligned} r &= y - \widehat{x} \\ &= y \circ k + \big(x - y\big) \circ k + \big(x - \widehat{x}\big) + \theta \\ &= y \circ k + n, \end{aligned} \tag{7.4}$$

where the small difference term $\big(x - y\big) \circ k$ has been included in a single noise term n, together with the denoising error $x - \widehat{x}$ and other disturbances. However, even in the case of perfect denoising, the noise term is likely to dominate the residual r which, therefore, will be only weakly correlated with the camera PRNU. In addition, denoising is typically less accurate in the presence of textured areas, and some signal components leak into the residual. This event lowers even dramatically the operative SNR, and makes the detection virtually impossible. Especially in these areas, the effectiveness of the denoising methods becomes crucial for the overall performance. To improve the residual estimation, the algorithm proposed in Section 7.3 relies on nonlocal denoising techniques.

In the following, we describe the image integrity verification procedure proposed in [41], which comprises three basic steps:

(i). estimation of the camera PRNU (off-line);

(ii). sliding-window pixel-wise forgery detection test;

(iii). morphological processing of test result map.

7.2.1 PRNU estimation

The first step of the procedure is not a hard task, because one is supposed to have either a large number of images taken by the camera of interest or the camera itself, and hence any reasonable estimation technique will provide, by sheer brute force, a good result. The maximum likelihood estimate of the PRNU from M given images $\{y_1, \dots, y_M\} \subset \mathbb{R}^N$ is computed in [41] as

$$(\forall \ell \in \{1, \dots, N\}) \qquad \widehat{k}^{(\ell)} = \frac{\sum_{m=1}^{M} y_m^{(\ell)} \, r_m^{(\ell)}}{\sum_{m=1}^{M} \left(y_m^{(\ell)}\right)^2}, \tag{7.5}$$

where r_m denotes the noise residual of the image y_m. Note that the weighting terms y_m account for the fact that dark areas of the image present an attenuated PRNU, and hence should contribute less to the overall estimate. In the following, for the sake of simplicity, we neglect the estimation error and assume to know the camera PRNU perfectly, that is $\widehat{k} = k$.

7.2.2 Forgery detection test

The detection problem can be formulated as a binary test between the hypothesis H_0 that the camera PRNU is absent (i.e., the pixel has been tampered) and the hypothesis H_1 that it is present (i.e., the pixel is genuine), yielding

$$(\forall \ell \in \{1, \dots, N\}) \qquad \begin{cases} H_0: & r^{(\ell)} = n^{(\ell)}, \\ H_1: & r^{(\ell)} = y^{(\ell)} k^{(\ell)} + n^{(\ell)}. \end{cases} \tag{7.6}$$

The true and estimated pixel classes will be denoted by binary values $u^{(\ell)}$ and $\widehat{u}^{(\ell)}$, both defined in $\{0, 1\}$. Notice that, since we focus on the detection of forgeries (signaled by the *absence* of the PRNU), the role of the two hypotheses is inverted w.r.t. what is usual. For example, we will talk of *False Alarm* when H_0 is accepted but H_1 holds (a forged pixel is declared genuine), and of *Missing Detection* when H_0 is accepted but H_1 holds (a genuine pixel is declared forged).

The detection test is based on the normalized correlation [41] between the restrictions of r and $y \circ k$ to a window $\mathcal{W}_\ell \subset \{1, \dots, N\}$ centered at ℓ, that is

$$\rho^{(\ell)} = \operatorname{corr}\left(\left[r^{(n)}\right]_{n \in \mathcal{W}_\ell}, \left[y^{(n)} k^{(n)}\right]_{n \in \mathcal{W}_\ell}\right), \tag{7.7}$$

where the normalized correlation of two vectors a and b is defined as

$$\operatorname{corr}(a, b) = \frac{(a - \overline{a})^\top (b - \overline{b})}{\|a - \overline{a}\| \, \|b - \overline{b}\|}, \tag{7.8}$$

with $\overline{a}$ (resp. $\overline{b}$) being a vector containing the arithmetic mean of a (resp. b). For each site $\ell \in \{1, \dots, N\}$, the pixel label $\widehat{u}^{(\ell)}$ is obtained by comparing the decision statistic $\rho^{(\ell)}$ with a threshold $\gamma_1 \in \mathbb{R}$, leading to the binary rule

$$\widehat{u}^{(\ell)} = \begin{cases} 0, & \text{if } \rho^{(\ell)} < \gamma_1, \\ 1, & \text{otherwise.} \end{cases} \tag{7.9}$$

The above threshold is selected with a Neyman-Pearson approach to obtain the desired false acceptance rate (FAR), namely a suitably small probability that a tampered pixel be wrongly identified as genuine. The probability distribution of ρ under hypothesis H_0 is estimated by computing the correlation between the camera PRNU and a large amount of noise residuals coming from other cameras, and using standard density fitting techniques. Rather large square blocks must be considered in order to obtain a reliable statistic, such as 128×128 pixels [41].

Even in the absence of forgery, the correlation might happen to be very low when the image is dark (since y multiplies the PRNU), saturated (because of intensity clipping), or when denoising does not perform well and some image content leaks into the noise residual. In [41], this problem is addressed by means of a "predictor" which computes the expected value $\widehat{\rho}^{(\ell)}$ of the correlation index under the hypothesis H_1, relying on some local images features such as texture, flatness and intensity. The principle consists of labeling the pixel $y^{(\ell)}$ as genuine (the less risky decision) when the prediction $\widehat{\rho}^{(\ell)}$ is too low, irrespective of the value of $\rho^{(\ell)}$. In this case, in fact, even for a genuine pixel, one cannot expect a correlation index much larger than 0. Therefore, the test becomes

$$\widehat{u}^{(\ell)} = \begin{cases} 0, & \text{if } \rho^{(\ell)} < \gamma_1 \text{ and } \widehat{\rho}^{(\ell)} > \gamma_2, \\ 1, & \text{otherwise.} \end{cases} \tag{7.10}$$

The second threshold $\gamma_2 \in \mathbb{R}$ is chosen heuristically by the user and separates, in practice, reliable regions from problematic ones. It is worth underlining that the refined decision test (7.10) can only reduce the false alarm rate but does not increase (actually it might reduce) the probability of detecting an actual forgery. In addition, the choice of the threshold itself is not obvious and can significantly impact on the performance. For this reason, the algorithm proposed in Section 7.3 employs a more flexible Bayesian approach that allows us to make decisions jointly on the whole image, rather than individually for each pixel.

7.2.3 Label map post-processing

The output from the previous step, being generated by decisions taken independently for each pixel, results in a fragmented and inconsistent map $\widehat{u}$ that needs to be post-processed. To do so, an *ad-hoc* morphological filtering is performed on the regions of pixels declared as forged, over a background of genuine pixels. In [41], all regions smaller than 64×64 pixels (one fourth of the window size) are attributed to random errors and removed. Finally, the surviving regions are dilated with a structured element of radius 20 pixel to approximately restore the shape of the forged region, since many border points go lost because their correlation index is computed on mixed (forged/genuine) blocks.

An important limitation of the reference method [41] is that the spatial dependencies exhibited by natural images are taken into account in the post-processing phase, after that the decisions about forged pixels have been already taken. On the contrary, the algorithm proposed in Section 7.3 aims at solving a binary segmentation problem involving a regularity constraint, so as to provide a smooth output by penalizing maps with isolated points or many small regions.

7.3 Proposed algorithm

The proposed algorithm relies on [41], using the same hypotheses, basic approach and tools, but differs profoundly from it in the formulation of the problem and, consequently, in the solution techniques. We follow a Bayesian approach which allows us to better balance the observed statistics and the prior knowledge on the image. As a consequence, we obtain an improved performance, namely an increased probability to reveal small forgeries, and a much lower probability of declaring forgery in genuine regions. In addition, as proposed in [47], we replace the wavelet-based filter used in [41] with a better nonlocal filter, which provides us with more reliable data for the subsequent decision phase.

7.3.1 Bayesian Formulation

Our goal is to find the label map $\widehat{u} \in \{0,1\}^N$ which has the maximum probability to occur given the observed image y, leading to

$$\widehat{u} = \underset{u \in \{0,1\}^N}{\arg\max} \ \Pr(u|y). \tag{7.11}$$

All quantities of interest, except for the PRNU, are modeled as random fields, using the same symbol Pr *for denoting their probability distributions.*

Like in [41], however, we consider the noise residual $r = y - f(y)$ in place of the original image, because of its reduced noise power, although this is sub-optimal in principle, because r is not a sufficient statistic for our decision problem [132].

We compute the decision statistics ρ from r. Under the hypothesis H_0 (forged pixel), the expected value of $\rho^{(\ell)}$ is zero, since the noise is assumed to be signal-independent. Under the hypothesis H_1, instead, the expected value of $\rho^{(\ell)}$ is larger than zero, but not known, as it depends in a complex way on local image features. Since this information is necessary to make any sensible decision, we resort to a further statistic, the predictor $\widehat{\rho}^{(\ell)}$ proposed in [41], and assume

$$(\forall \ell \in \{1, \dots, N\}) \qquad \widehat{\rho}^{(\ell)} \simeq \mathbb{E}\{\rho^{(\ell)} | H_1\}. \tag{7.12}$$

Therefore, according to the conditional Bayes law, our problem becomes

$$\begin{aligned} \widehat{u} &= \underset{u \in \{0,1\}^N}{\arg\max} \ \Pr(u|\rho, \widehat{\rho}) \\ &= \underset{u \in \{0,1\}^N}{\arg\max} \ \Pr(\rho|u, \widehat{\rho}) \Pr(u|\widehat{\rho}) \\ &= \underset{u \in \{0,1\}^N}{\arg\max} \ \Pr(\rho|u, \widehat{\rho}) \Pr(u), \end{aligned} \tag{7.13}$$

in which the last equality comes from the fact that $\widehat{\rho}$ does not depend on u, but only on the image content, be it genuine or forged. Hereabove, the term $\Pr(\rho|u, \widehat{\rho})$ is the conditional likelihood of observing ρ, while $\Pr(u)$ accounts for the prior probability of the labels.

The model in (7.13) provides some insight about the strength of the Bayesian approach. The prior term, in fact, allows us to take into account all available knowledge on the expected forgery map, so as to guide the decision process towards reasonable results. In the absence of such a term, decisions are typically taken independently for each pixel, which could generate fragmented and inconsistent maps, calling for intense *ad hoc* postprocessing. On the contrary, by

choosing a model for u that exploits the strong spatial dependencies exhibited by natural images, decisions are taken jointly, penalizing ultimately maps with isolated points or many small regions and providing a smooth output.

The choice of the prior model plays a pivotal role for performance and, not least, for the complexity of the optimization algorithm, which is why we use a Markovian prior. Markov random fields (MRF) [16, 83, 103, 149] are relatively simple and effective tools to model the prior distribution of an image. An image $u \in \mathbb{R}^N$ is said to be a MRF when each pixel $u^{(\ell)}$ depends on the rest of the image only through a selected group of neighbors $\mathcal{N}_\ell \subset \{1,\dots,N\} \setminus \{\ell\}$, namely

$$\Pr\Big(u^{(\ell)} \mid \big(u^{(n)}\big)_{n\in\{1,\dots,N\}\setminus\{\ell\}}\Big) = \Pr\Big(u^{(\ell)} \mid \big(u^{(n)}\big)_{n\in\mathcal{N}_\ell}\Big). \tag{7.14}$$

Thanks to the Markovian property, one avoids the challenging problem of assigning a global prior, and specifies only statistics on the local neighborhoods. It can be proved that any positive MRF has a Gibbs probability law defined as

$$\Pr(u) = \frac{1}{Z}\exp\Big\{-\sum_{e\in\mathcal{C}} U_e(u)\Big\}, \tag{7.15}$$

where Z is a normalizing constant, and C contains a number of subset of $\{1,\dots,N\}$ called *cliques*. The potentials U_e are ultimately responsible for the MRF properties. Consider for example a two-pixel clique $e = (\ell, n)$, with associated potential $U_{\ell,n}(u) = \lambda|u^{(\ell)} - u^{(n)}|$ for some $\lambda > 0$. If $u^{(\ell)} = u^{(n)}$, the clique will contribute a unitary factor to the overall probability of the image, while if $\lambda|u^{(\ell)} - u^{(n)}| = \Delta \gg 0$ it will contribute a factor $e^{-\Delta}$ very close to 0. With this choice of local potentials, images with a sharp transition between pixels ℓ and n are made very unlikely *a priori*, and smooth images are preferred.

In our case, since $u \in \{0,1\}^N$, we resort to the popular Ising model [149], where only single-site and two-site cliques are considered, yielding the potentials

$$(\forall \ell \in \{1,\dots,N\}) \qquad U_\ell(u) = \alpha\big(u^{(\ell)} - 1/2\big) = \begin{cases} -\alpha/2, & \text{if } u^{(\ell)} = 0, \\ +\alpha/2, & \text{if } u^{(\ell)} = 1, \end{cases} \tag{7.16}$$

and, for every $\ell \in \{1,\dots,N\}$,

$$(\forall n \in \mathcal{N}_\ell) \qquad U_{\ell,n}(u) = \lambda|u^{(\ell)} - u^{(n)}| = \begin{cases} \lambda, & \text{if } u^{(\ell)} \neq u^{(n)}, \\ 0, & \text{otherwise,} \end{cases} \tag{7.17}$$

where $\mathcal{N}_\ell$ denotes the set of 4-connected neighbors of ℓ. Single-site potentials are directly related to the prior probability p_0 and p_1 of the classes, as

$$\alpha = \log\big(p_0/p_1\big). \tag{7.18}$$

Two-site potentials, instead, penalize label transitions (remember the minus sign before the energy) between 4-connected sites, enforcing a bias toward smooth images whose strength depends on the edge-penalty parameter λ. For $\lambda = 0$, there is no interaction between pixels, no matter how close they are, while the bias against label transition grows stronger with increasing values of λ, and extends well beyond a local clique thanks to chain propagation. Therefore, in the absence of observable data, the prior probability is maximized by a flat map, with all labels equal to the one most probable *a priori*.

Turning to the likelihood term, we assume the conditional independence of ρ (this and other hypotheses will be discussed later). As a result, after taking the negative log, we can rewrite Problem (7.13) as

$$\underset{u\in\{0,1\}^N}{\text{minimize}} \quad \sum_{\ell=1}^{N} -\log \Pr\left(\rho^{(\ell)}|u^{(\ell)},\widehat{\rho}^{(\ell)}\right) + \sum_{\ell=1}^{N} U_\ell(u) + \sum_{\ell=1}^{N}\sum_{n\in\mathcal{N}_\ell} U_{\ell,n}(u), \quad (7.19)$$

which, after having set $f_\ell(u^{(\ell)}) = -\log \Pr\left(\rho^{(\ell)}|u^{(\ell)},\widehat{\rho}^{(\ell)}\right)$, leads to

$$\underset{u\in\{0,1\}^N}{\text{minimize}} \quad \sum_{\ell=1}^{N} f_\ell(u^{(\ell)}) + \alpha\sum_{\ell=1}^{N} u^{(\ell)} + \lambda\sum_{\ell=1}^{N}\sum_{n\in\mathcal{N}_\ell} |u^{(\ell)} - u^{(n)}|. \quad (7.20)$$

Since the function $f_\ell(u^{(\ell)})$ of a binary variable $u^{(\ell)} \in \{0,1\}$ can be expressed as $f_\ell\big(u^{(\ell)}\big) = u^{(\ell)}\big[f_\ell(1) - f_\ell(0)\big] + f_\ell(0)$, Problem (7.20) reduces to

$$\underset{u\in\{0,1\}^N}{\text{minimize}} \quad \sum_{\ell=1}^{N} c^{(\ell)}\, u^{(\ell)} + \lambda\sum_{\ell=1}^{N}\sum_{n\in\mathcal{N}_\ell} |u^{(\ell)} - u^{(n)}|, \quad (7.21)$$

where $c^{(\ell)} = f_\ell(1) - f_\ell(0) + \alpha$. Finally, we assume the likelihood to be Gaussian under both hypotheses, with zero mean and variance σ_0^2 under hypothesis H_0, and mean $\widehat{\rho}^{(\ell)}$ and variance σ_1^2 under hypothesis H_1, namely

$$f_\ell(1) = \frac{\big(\rho^{(\ell)} - \widehat{\rho}^{(\ell)}\big)^2}{2\sigma_1^2} + \log\sigma_1 + \log\sqrt{2\pi}, \quad (7.22)$$

$$f_\ell(0) = \frac{\big(\rho^{(\ell)}\big)^2}{2\sigma_0^2} + \log\sigma_0 + \log\sqrt{2\pi}, \quad (7.23)$$

and by recalling α in (7.18), we readily obtain the expression

$$c^{(\ell)} = \frac{\big(\rho^{(\ell)} - \widehat{\rho}^{(\ell)}\big)^2}{2\sigma_1^2} - \frac{\big(\rho^{(\ell)}\big)^2}{2\sigma_0^2} + \log\frac{p_0\,\sigma_1}{p_1\,\sigma_0}. \quad (7.24)$$

Problem (7.21) is the final expression to consider for finding the optimal $\widehat{u}$. Before that, however, let us gain some insight into the meaning of the objective function in (7.21). Assume $\lambda = 0$, for the time being, which means that decisions are taken independently for each pixel, selecting $u^{(\ell)} = 1$ if the biased likelihood $c^{(\ell)}$ is negative, and $u^{(\ell)} = 0$ otherwise. If we assume also $\sigma_0^2 = \sigma_1^2$ and $p_0 = p_1$ (no bias), the algorithm reduces to comparing the correlation index $\rho^{(\ell)}$ with a threshold $\widehat{\rho}^{(\ell)}/2$ placed halfway between the two means. Unequal prior probabilities and variances modify somewhat the decision regions, favoring for example the most probable class, without altering the essence of the procedure.

If we now consider $\lambda > 0$, decisions are not independent anymore, as transitions are penalized, and a decision can be reverted if convenient. This can happen especially if the biased likelihood $c^{(\ell)}$ is small, that is typically when $\widehat{\rho}^{(\ell)}$ is small. However, if the biased likelihood $c^{(\ell)}$ keeps the same sign over a relatively large and compact area, there is no reason to change decisions, or refrain from taking a decision at all, even if absolute values are small. Therefore, it becomes possible to detect relatively small forgeries even in dark and textured

regions of the image. This is the fundamental improvement w.r.t. the original algorithm, which did not trust at all the decision statistic in problematic regions. Both algorithms enforce some compactness constraints, but in [41] this is done only by morphological filtering of the label map, after irreversible hard decisions have been already taken. Here, on the contrary, likelihood and prior are weighted optimally (in the limit of the accuracy of the model) before taking any decision, and the compactness is taken into account by the $\boldsymbol{\ell}_1$-norm of pixel differences.

A few words are due about the hypotheses. The conditional independence of the likelihood does not hold true if correlation indexes are computed on largely overlapping sliding windows. A necessary condition to restore it is to use disjoint blocks, but this would entail an annoying loss of spatial resolution. Therefore, as a reasonable compromise, we use a subsampling of 8×8 (only in the optimization phase), which guarantees a good spatial resolution, allows for a significantly saving in CPU time, and reduces the spatial correlation of the likelihood field. On the other hand, the residual correlation, which modifies the absolute values of the global likelihood, is automatically taken into account through the edge penalty parameter λ, set by means of preliminary experiments.

Concerning the Gaussian model, it was observed in [41] that it fits experimental data very accurately under H_0 but not under H_1, where a generalized Gaussian (GG) model was preferred. However, the choice of the GG model is strongly influenced by a small number of outliers that lie on the right tail of the distribution (large $\rho^{(\ell)}$) of little interest for the decision. In the region where a good modeling is more important, between 0 and $\widehat{\rho}^{(\ell)}$, the simpler Gaussian fitting seems accurate enough. The ratio between the prior probabilities p_0/p_1 can be based on actual observations or, more practically, based on the different risks associated with false alarm and missing detection errors. This choice can be thus loosely related with the choice of the constant FAR threshold of [41].

7.3.2 Optimization by convex relaxation

Problem (7.21) can be solved, either exactly or approximatively, by a number of combinatorial algorithms, such as simulated annealing [134], iterated conditional modes [16], graph cuts [25], or belief propagation [230]. A powerful alternative approach relies on the concept of convex relaxation, which consists of converting Problem (7.21) into a convex one admitting the same solution [38]. This approach amounts to replace the discrete domain $\{0,1\}^N$ with the unitary hypercube $[0,1]^N$, obtaining eventually the convex optimization problem

$$\underset{u\in[0,1]^N}{\text{minimize}} \quad \sum_{\ell=1}^{N} c^{(\ell)}\, u^{(\ell)} + \lambda \sum_{\ell=1}^{N} \sum_{n\in\mathcal{N}_\ell} |u^{(\ell)} - u^{(n)}|. \tag{7.25}$$

It can be proved [38] that (7.25) is equivalent to (7.21), in the sense that a solution to the integer-valued problem can be obtained by thresholding at any level $\mu \in]0,1[$ the solution to the above convex problem.

The great advantage of convex relaxation is that Problem (7.25) can be equivalently rewritten in the following constrained form

$$\underset{u\in[0,1]^N}{\text{minimize}} \quad \sum_{\ell=1}^{N} c^{(\ell)}\, u^{(\ell)} \quad \text{s.t.} \quad \sum_{\ell=1}^{N} \sum_{n\in\mathcal{N}_\ell} |u^{(\ell)} - u^{(n)}| \leq \eta, \tag{7.26}$$

as these two formulations are equivalent for specific values of λ and η. As already argued in the previous chapters, Problem (7.26) may be more practical, as the choice of η can be related to some physical properties of the target signal [61]. This is especially true in binary segmentation, where the edge penalty is directly related to the forgery size, and thus the upper bound η can be reasonably set according to the outcome of a number of controlled preliminary experiments.

The standard way of solving Problem (7.26) via proximal methods leads to a sequence of steps involving the projection onto the $\boldsymbol{\ell}_1$-ball. While this projection can be computed through specific numerical procedures [71, 217], a more efficient approach consists of resorting to the proposed epigraphical splitting. The latter amounts to reformulating Problem (7.26) as

$$\underset{(u,\zeta)\in[0,1]^N\times V}{\text{minimize}} \quad c^\top u \quad \text{s.t.} \quad (Fx,\zeta)\in E, \tag{7.27}$$

where $c=(c^{(\ell)})_{1\leq\ell\leq N}$, F is defined in (3.7) with $\omega_{1,n}\equiv 1$ and $M_\ell\equiv 4$, while

$$\begin{aligned} E &= \big\{(y,\zeta)\in\mathbb{R}^{4N}\times\mathbb{R}^{4N} \;\big|\; (\forall \ell\in\{1,\dots,4N\}) \quad (y^{(\ell)},\zeta^{(\ell)})\in \operatorname{epi}|\cdot|\big\} \\ V &= \big\{\zeta\in\mathbb{R}^{4N} \;\big|\; 1_{4N}^\top\zeta\leq\eta\big\}. \end{aligned} \tag{7.28}$$

The iterations of M+LFBF associated to the above problem are reported in Algorithm 7.1, where the projection P_E is computed by virtue of Proposition (2.3.2). The algorithm stops when either the convergence criterion

$$\|u^{(i+1)}-u^{(i)}\|<10^{-5}\|u^{(i)}\| \tag{7.29}$$

is satisfied or the maximum allowed number of iterations is reached. Experiments show that 1000 iterations are always enough for a good convergence.

Algorithm 7.1 *Iterations of M+LFBF [66] for Problem* (7.27)

INITIALIZATION

$$\left\lfloor\begin{array}{l} \text{choose } \big(x^{[0]},\zeta^{[0]}\big)\in\mathbb{R}^N\times\mathbb{R}^{4N} \\ \text{choose } \big(y^{[0]},\xi^{[0]}\big)\in\mathbb{R}^{4N}\times\mathbb{R}^{4N} \\ \text{set } \gamma\in\,]0,1/4[\end{array}\right.$$

FOR $i=0,1,\dots$

$$\left\lfloor\begin{array}{l} p^{[i]}=P_{[0,1]^N}\Big(x^{[i]}-\gamma\,(c+F^\top y^{[i]})\Big) \\ \rho^{[i]}=P_V\big(\zeta^{[i]}-\gamma\,\xi^{[i]}\big) \\ \widehat{y}^{[i]}=y^{[i]}+\gamma F x^{[i]} \\ \widehat{\xi}^{[i]}=\xi^{[i]}+\gamma\,\zeta^{[i]} \\ (v^{[i]},\nu^{[i]})=(\widehat{y}^{[i]},\widehat{\xi}^{[i]})-\gamma P_E\Big(\widehat{y}^{[i]}/\gamma,\widehat{\xi}^{[i]}/\gamma\Big) \\ y^{[i+1]}=v^{[i]}+\gamma F\big(p^{[i]}-x^{[i]}\big) \\ \xi^{[i+1]}=\nu^{[i]}+\gamma\big(\rho^{[i]}-\zeta^{[i]}\big) \\ x^{[i+1]}=p^{[i]}-\gamma F^\top(v^{[i]}-y^{[i]}) \\ \zeta^{[i+1]}=\rho^{[i]}-\gamma\big(\nu^{[i]}-\xi^{[i]}\big) \end{array}\right.$$

7.3.3 Nonlocal denoising

Although the proposed Bayesian approach will very likely improve the reliability of forgery detection, performance ultimately depends mostly on the quality of the original data. Looking at the definition of noise residual in (7.4), we see that the local signal-to-noise ratio depends primarily on image intensity y (which multiplies the PRNU) and the strength of three noise terms. Among them, the denoising error $x - \widehat{x}$ typically predominates due to imperfect treatment of edges and textures, and thus the quality of denoising impacts immediately and strongly on the detection performance.

Most denoising algorithms separate signal from noise based on their spectral properties, assuming the signal to be dominant at the lower spatial frequencies, and noise to prevail at the higher ones. This approach inspires both spatial-domain and transform-domain filters, in more or less explicit forms, including those operating in the wavelet domain. Unfortunately, such a spectral separation holds very well in flat areas of the image, but much less so in textured areas, where the signal has significant components at the higher frequencies of the spectrum. As a result, some signal components are treated as noise and filtered, causing smoothing and blurring in the output image and, what is worse in our point of view, contributing significantly to noise power in the residual.

To improve the quality of noise residual, we resort to nonlocal denoising [32], and in particular to the state-of-the-art BM3D algorithm [79]. As the name suggests, nonlocal filtering estimates the true value of a given pixel by relying on pixels (taken anywhere in the image) statistically homogeneous with the target. Thanks to the inherent self-similarity of images, there are several patches in the surroundings of the target that happen to be very similar to it, and therefore can be assumed to be good predictors. By jointly filtering such patches, one obtains a good estimate of the target, with little or no spatial smearing, mimicking a true statistical average. The joint processing of similar patches ensures that the original image is well estimated also in the presence of strong texture. As a result, the noise residual obtained using BM3D presents limited traces of signal, leading to a better correlation index field, especially in the textured areas.

7.4 Experimental results

In this section, we assess experimentally the performance of the proposed technique. Experiments are carried out on four cameras: Canon EOS 450D, Canon IXUS 95IS, Nikon D200, and Nikon Coolpix S5100. For each camera, we use a first training set of 200 images to estimate the PRNU pattern, and another training set of 20 images to design the predictor. Performance indicators are then computed on a larger set, disjoint from the training sets, comprising 600 images for each camera. All training and test images have the same size of 768×1024 pixels, and are cropped from the same region of the original images output by the camera. For each test image, we consider both the genuine version, used to look for false alarms (wrongly declared forgeries), and a forged version, used to look for correct detections (correctly declared forgeries). To create the forged version, we replace a square at the center of the image with an equal-size square taken at random (but in a reproducible way) from another image of the same

(a) Genuine image.

(b) Tampered image.

(c) Decision map for the genuine image.

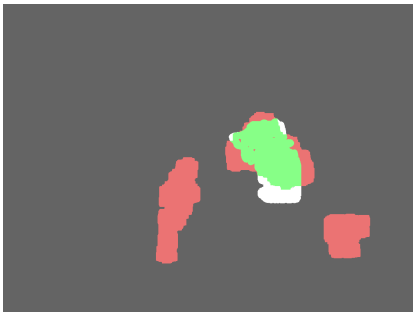

(d) Decision map for the tampered image.

Figure 7.1 *Color coded detection masks. Gray: genuine pixel declared genuine; red: genuine pixel declared tampered (error); white: tampered pixel declared genuine (error); green: tampered pixel declared tampered.*

or different camera. To study how performance depends on forgery size, we use forgeries of three sizes, 128×128, 256×256, and 384×384 pixels, creating thus three test subsets of 200 images each.

In order to assess separately the improvements due to the Bayesian formulation and the improved filtering, we consider all four combinations in our experiments, that is, the original technique with constant false acceptance rate (CFAR) decision rule and either Mihcak or BM3D denoising, and the proposed version with the Bayesian decision rule and, again, Mihcak or BM3D denoising. In the following, we will call these techniques for short CFAR-Mihcak, CFAR-BM3D, Bayes-Mihcak and Bayes-BM3D, respectively.

Figure 7.1 shows an example genuine image together with a tampered version with a medium-size realistic forgery at the center and, on the bottom row, the corresponding color-coded decision masks output by the CFAR-Mihcak algorithm. The first mask, obtained for the genuine image, is used to compute the false alarm probability, as the ratio between the number of pixels declared forged (in red) and the image size. The second mask, obtained for the forged image, is used to compute the detection probability, as the ratio between the number of correctly identified forged pixels (in green) and the forgery size (white + green).

7.4.1 Pixel-level analysis

Figure 7.2 shows, for each of the four cameras, the receiver operating characteristics (ROCs) of the four algorithms, computed on the complete test set (600 images with forgeries of various sizes). To improve the readability, we show only a close-up of the most relevant top-left region. Each ROC is computed by varying in a wide range the algorithm main parameters, γ_1 and γ_2 for CFAR, α and η for Bayes, and then taking the upper envelope of the resulting points (P_{FA}, P_D). Although CFAR-Mihcak can be considered one of the best forgery detection techniques known to date, Bayes-BM3D technique improves clearly upon it, showing for all cameras a uniformly better ROC. Both the improved denoising filter and the global Bayesian formulation provide significant improvements over CFAR-Mihcak. Their joint use, however, improves performance still further, showing that better data do not solve all problems by themselves, nor is sufficient to adopt a more clever decision strategy irrespective of data quality. Of course, results change slightly from camera to camera but the general behavior is always the same. For the Canon IXUS, the ROCs are much closer to one another. For the two Nikon cameras, the performance is generally worse than with the Canon cameras. This is not surprising since the performance depends strongly on the average intensity of the PRNU noise, which varies significantly for different manufacturers and camera models, and is somewhat smaller for the Nikon. This fact explains also why, for the Nikon cameras, unlike the Canon, CFAR-BM3D outperforms Bayes-Mihcak. In fact, given the weaker PRNU, improving data reliability is more rewarding than using a better optimization strategy on unreliable data.

To gain insight into how the performance depends on forgery size, Figure 7.3 shows in distinct graphs the ROCs (same close-up as before) computed only on large (384×384 pixels), medium (256×256), and small size (128×128) forgeries. To save space, results are shown only for the Canon EOS camera, but the same general behavior is observed in all cases. The performance is always very good (remember that we are zooming on the upper-left part of the (P_{FA}, P_D) domain) but, as expected, gets worse and worse as the forgery size decreases, because some small forgeries might be missed altogether. Note that the gap between the Bayesian and CFAR approach grows very wide in the case of large forgeries, where the former exhibits a near-optimal behavior. In fact, the Bayesian approach allows one to detect large forgeries even in unfavourable conditions, such as dark and textured regions, where CFAR approach can be fooled.

Although the *pixel-level* false alarm and detection probabilities reported in the above figures are widespread performance indicators in this field, they are not fully appropriate to assess forgery detection performance. A typical user is mostly interested, in order of decreasing importance, in: (i) establishing whether an image is tampered, (ii) finding the approximate location of detected forgeries, (iii) knowing their approximate size and shape. In fact, once a forgery has been detected, together with its approximate location, one can resort to many other tools to obtain more detailed information. Moreover, automatic forgery detection can be used to pre-screen a large number of images, in order to select those more likely to have been tampered with for visual analysis by expert photo-interpreters. Therefore, in the following, we present two more sets of curves, the customary *image-level* results, and what we call *object-level* results, concerning the probability of correctly detecting forged objects within an image.

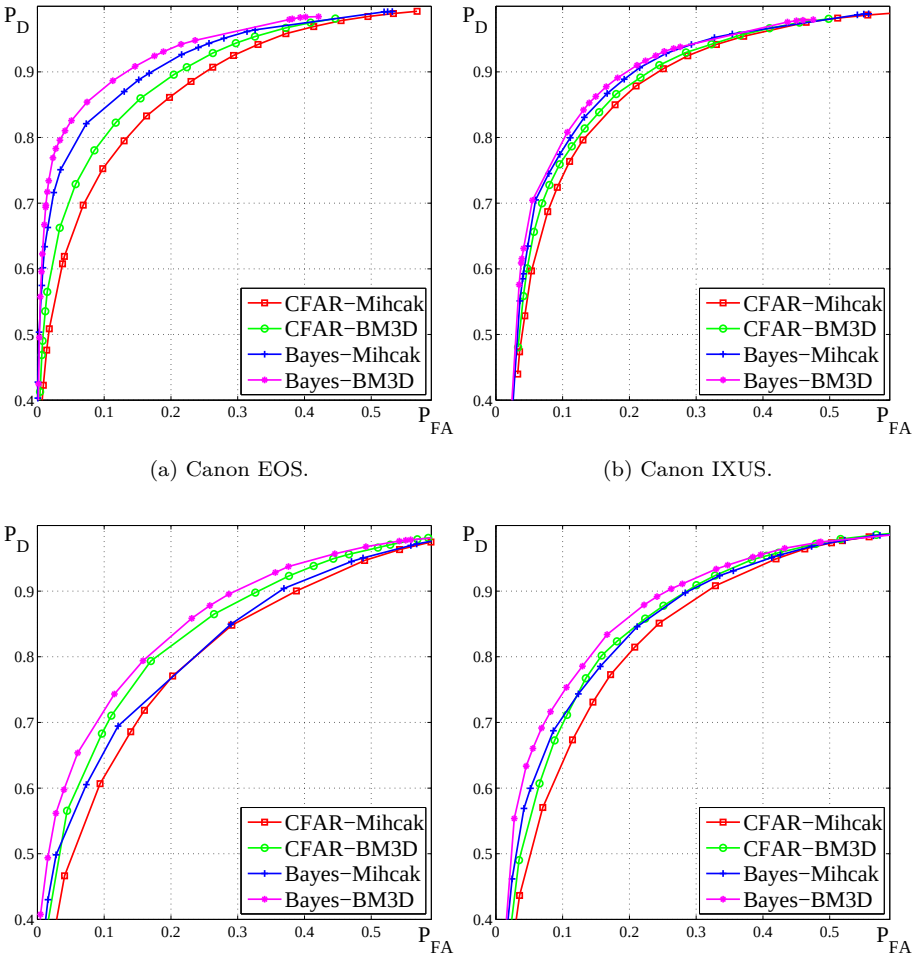

(a) Canon EOS.

(b) Canon IXUS.

(c) Nikon D200.

(d) Nikon Coolpix.

Figure 7.2 *Pixel-level ROCs (close-up) for all cameras under test.*

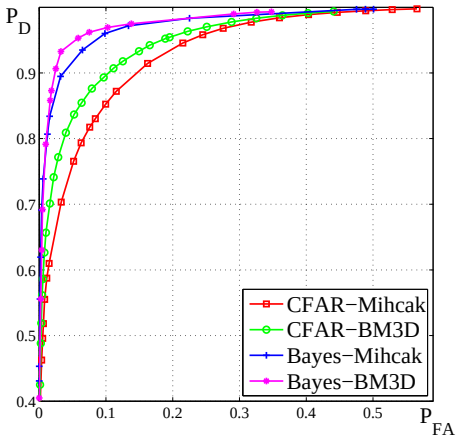

(a) Large-size forgeries.

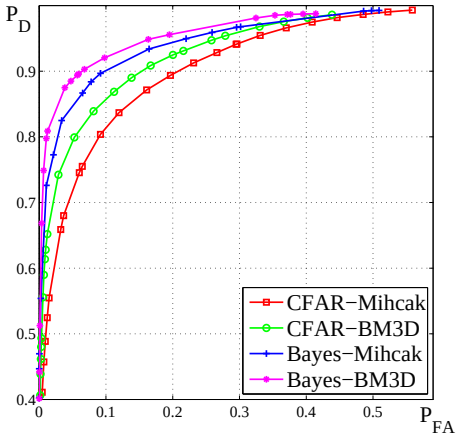

(b) Medium-size forgeries.

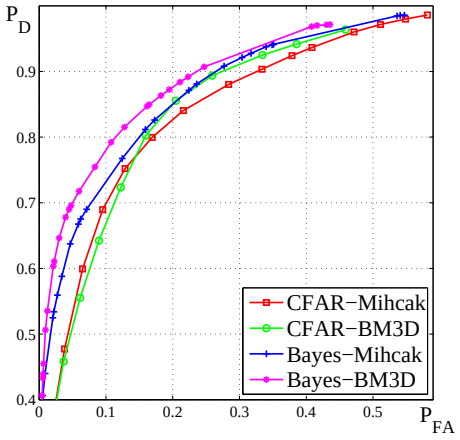

(c) Small-size forgeries.

Figure 7.3 *Pixel-level ROCs (close-up) for the Canon EOS camera.*

7.4.2 Image-level analysis

To compute image-level results, we consider only the global decision made on the whole image. More precisely, the image under test is declared forged if any of its pixels is (remember that small regions are erased right away from the map), and genuine otherwise. False alarm and detection probabilities are then computed as the fraction of genuine and, respectively, forged images that are declared forged by the algorithm. Figure 7.4 shows the results for the four cameras under test, computed over the complete test set. The general behavior is the same as in the case of pixel-level ROCs, with the only difference that all curves are now further from the top-left corner (ideal performance). However, this is obvious considering that a wrong decision on just a small number of pixels may cause a wrong image-level decision on the whole image.

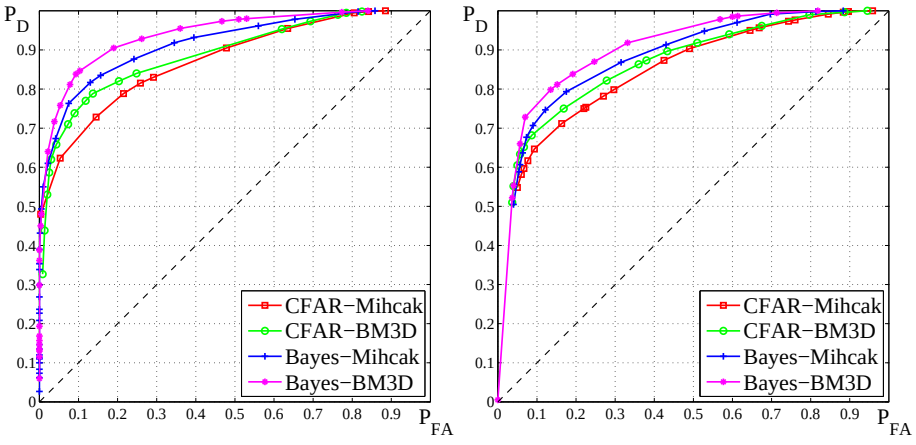

(a) Canon EOS. (b) Canon IXUS.

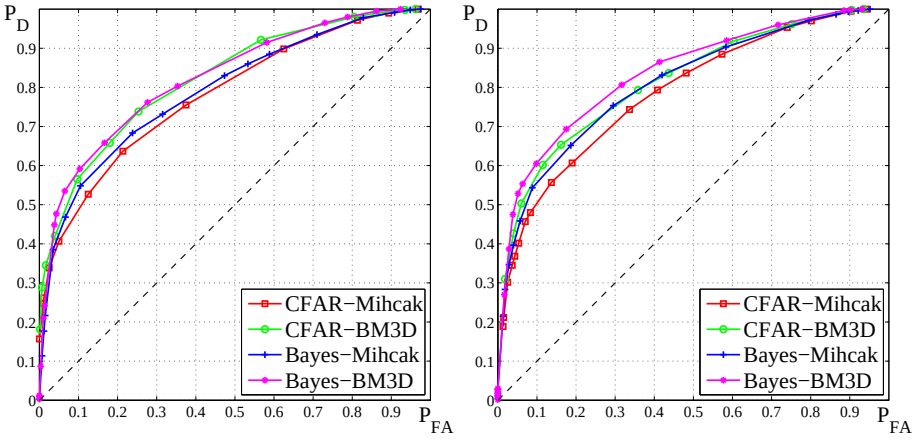

(c) Nikon D200. (d) Nikon Coolpix.

Figure 7.4 *Image-level ROCs for all cameras under test.*

7.4.3 Object-level analysis

Image-level ROCs provide information on the ability of an algorithm to classify images as forged or genuine. A more ambitious goal, however, is to localize forgeries with a reasonable accuracy. The PRNU-based approach, unlike many others, has the potential to address this task, therefore it is certainly worth carrying out a specific experimental analysis. To this end, we compute object-level ROCs, by modifying the definition of correct detection w.r.t. the image-level case. More precisely, we declare a correct detection only when the binary mask output by the algorithm covers a significant fraction, $\omega \in [0, 1]$, of the actual forged region, namely, going back to the example of Figure 7.1, if the green area is large enough w.r.t. the forgery. Therefore, all situations in which the output map covers just a tiny part of the actual forgery, or is even disjoint with it, providing little or no hints for visual inspection, are regarded as missing detections.

Figure 7.5 reports, for all cameras, the object-level performance computed on the complete test set with $\omega = 0.2$. As expected, all ROCs drift downward w.r.t. the corresponding image-level curves. The impairment, however, is almost negligible for the Bayesian algorithms, with either Mihcak or BM3D denoising, while it is dramatic for the CFAR techniques, especially for the Nikon cameras. In hindsight, this could be expected, since the MRF prior drives the optimization towards a compact and well localized output mask, while no guarantees in this sense is given by the CFAR-based thresholding. It is also interesting that ROCs are now in the same order of performance for all cameras, Bayes-BM3D > Bayes-Mihcak > CFAR-BM3D > CFAR-Mihcak, confirming that at object-level the most important improvement comes from making decisions globally as opposed to locally. Figure 7.6 shows, with reference only to the Canon EOS, that the Bayes techniques have an object-level performance almost independent of ω, while for the CFAR techniques a clear dependence is observed, indicating a less accurate forgery localization.

We complete this analysis by reporting, in Figure 7.6, again for the Canon EOS camera, the object-level ROCs computed separately on large (384×384 pixels), medium (256×256) and small (128×128) forgeries. As expected, performance drops for all algorithms as the forgery size reduces, and it is quite poor for the smallest size, which coincides with the size of the window used to compute the decision statistics. In this condition, the statistics used to detect forged pixels are always (except for a single point) computed on mixed data, thus impairing their diagnostic power. This is certainly one of the major limitations of the PRNU approach, and probably the main topic to deal with in future research. Barring this limiting case, performance is generally good, especially for the proposed version of the technique, which provides always a significant gain w.r.t. the original one.

In summary, this performance analysis shows that the proposed technique is not only able to reliably classify a test image as genuine or forged, but also to localize with good accuracy the detected forgery, provided it is large enough w.r.t. the analysis window. Under all these respects, it improves significantly over the reference technique.

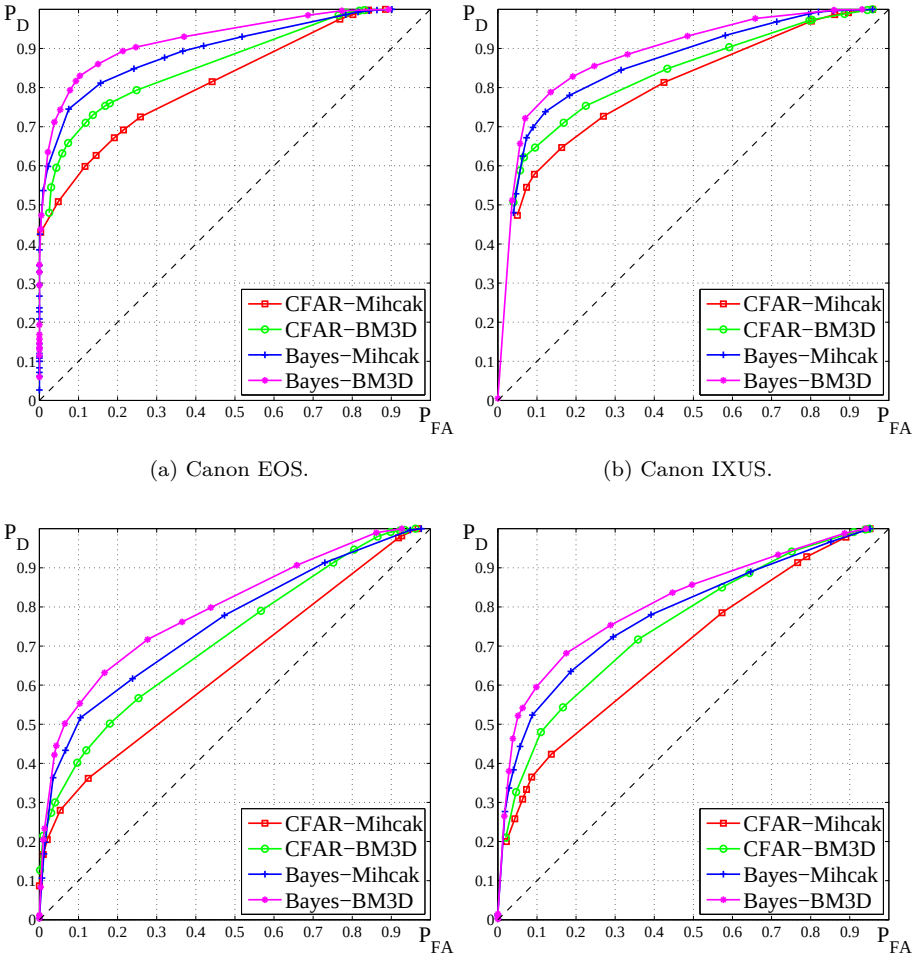

(a) Canon EOS.

(b) Canon IXUS.

(c) Nikon D200.

(d) Nikon Coolpix.

Figure 7.5 *Object-level ROCs ($\omega = 0.2$) for all cameras under test.*

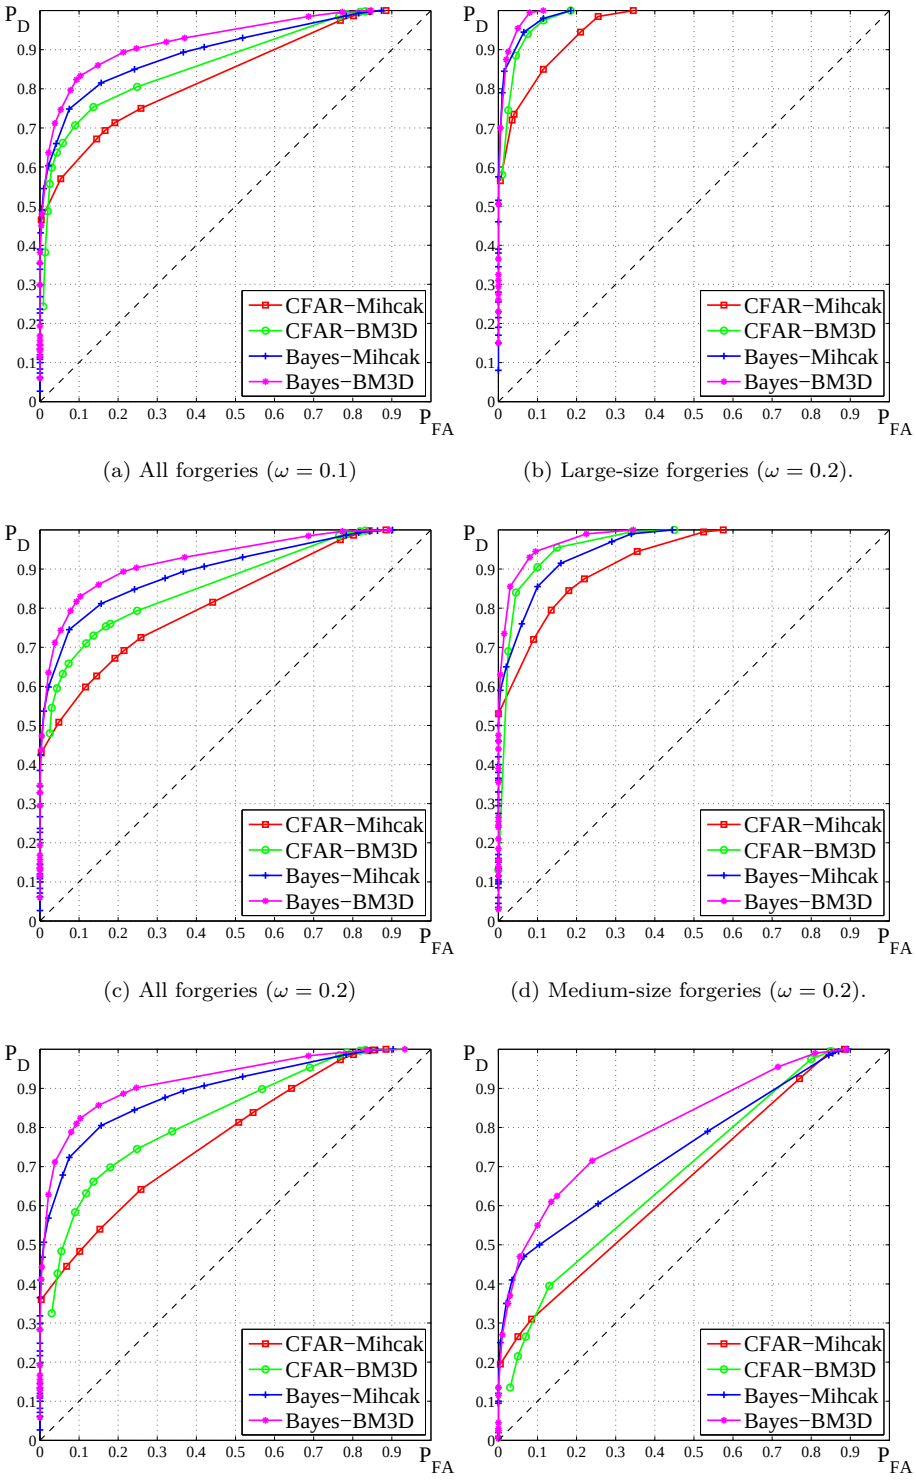

(a) All forgeries ($\omega = 0.1$)

(b) Large-size forgeries ($\omega = 0.2$).

(c) All forgeries ($\omega = 0.2$)

(d) Medium-size forgeries ($\omega = 0.2$).

(e) All forgeries ($\omega = 0.4$)

(f) Small-size forgeries ($\omega = 0.2$).

Figure 7.6 *Object-level ROCs for Canon EOS camera.*

7.4.4 Sample results on realistic cases

To gain a better insight into the behavior of proposed and reference algorithms, we now consider some sample results obtained by running the algorithms on a set of real-world tampered images taken by our four cameras. We consider both the insertion of a new object into the scene (splicing), and the cancellation of an object from the scene, realized by copy-moving part of the background to hide it. For both CFAR-Mihcak and Bayes-BM3D, we use the set of parameters corresponding to the best operating points in their respective object-level ROCs with $\omega = 0.20$. As an example, for the Canon EOS-450 camera, we use $\alpha = \log 3, \eta = 48$ for Bayes-BM3D, and $\gamma_1 = 0.026, \gamma_2 = 0.013$ for CFAR-Mihcak. Here, we focus on some situations where clear differences arise in the behavior of algorithms.

Figure 7.7 shows the outputs of CFAR-Mihcak and the Bayes-BM3D algorithms on some examples. The first one (Camel) corresponds to a simple case, with a large, bright and smooth forgery which, in fact, is successfully and accurately detected by both algorithms. Note that the correlation fields computed after Mihcak and BM3D filtering follow closely the predicted field, except in the region of the forgery, where the original PRNU is missing and a large deviation is observed. Notice also that no visual hint suggests the presence of a forgery, making the detection virtually impossible without proper tools.

The second example (Road Sign) is similar to the previous one, except for the two forgeries that are much smaller than before, about 128×128 pixels. As a consequence, the CFAR-Mihcak algorithm misses both of them, while Bayes-BM3D detects them correctly. Notice, in the center-bottom of the computed index correlation fields, the reddish square region with sharp edges caused by a single defective sensor, a feature that could be exploited to improve performance.

The third image (Wading Bird) is much darker than the previous ones and also more textured. As a consequence, the computed correlation indexes are quite small (blueish), especially after Mihcak filtering, which might easily induce false alarms. The optimum thresholds of the original algorithm, set so as to limit false alarms, discard the forgery region as well, causing a missed detection. The proposed technique, instead, detects accurately the spliced object, thanks to both the better filtering and the more sophisticated decision strategy.

The problem with textures is even more evident in the fourth example (Beach), where both algorithms detect easily the copy-move forgery, but CFAR-Mihcak presents also a false alarm caused by the bird feathers, highly textured, and imperfectly filtered by the wavelet-based denoiser. Also in the subsequent example (Roof), the texture causes some problems to CFAR-Mihcak. The algorithm detects a forged area which only partially (less than 10%) overlaps the actual spliced object and is therefore classified as a false alarm. This decision appears to be correct in this case, since the algorithm seems to follow the textured roof rather than the dark bird.

Finally, we show the case of a dark and textured image (Market) where both algorithms fail, detecting the forgery but generating multiple false alarms as well. With such a low quality image, on the other hand, it is difficult to envision successful strategies. We do not even show the case of forgeries much smaller than the analysis window, since they are obviously missed by definition by both algorithms. Notice, however, going back to the Wading Bird example, that the spliced object is indeed much thinner than 128 pixels, although longer than that, but still detected by the Bayesian strategy.

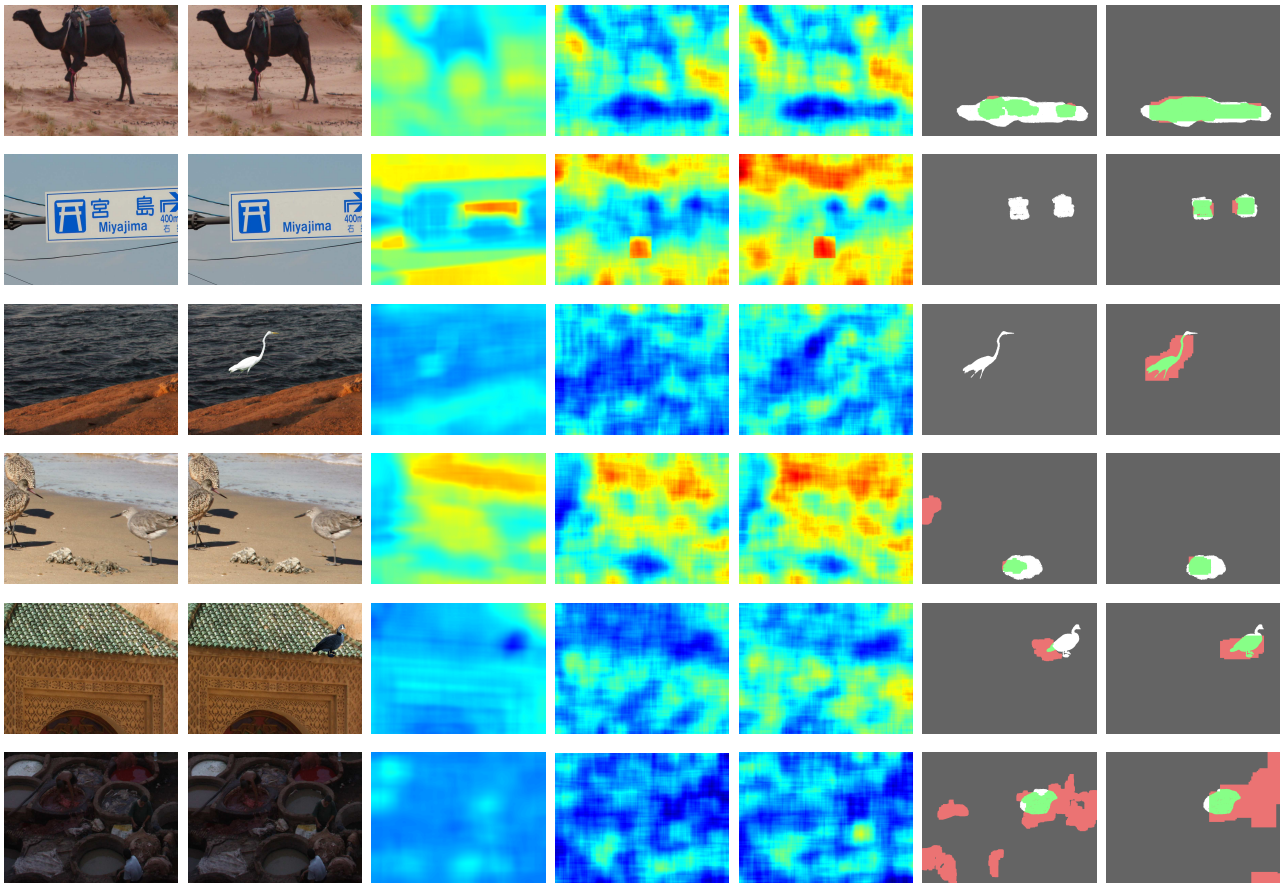

(a) (b) (c) (d) (e) (f) (g)

Figure 7.7 *Forgery detection results for some selected examples: Camel, Road Sign, Wading Bird, Beach, Roof, Market.*
(a) Original image.
(b) Tampered version.
(c) Predicted correlation field $\widehat{\rho}$.
(d) Correlation index field ρ computed after Mihcak filtering.
(e) Correlation index field ρ computed after BM3D filtering.
(f) Color-coded detection mask provided by CFAR-Mihcak.
(g) Color-coded detection mask provided by Bayes-BM3D.

7.4.5 Comparison of running times

Table 7.1 reports the average CPU times spent on a 3.40Ghz desktop computer with 8GB memory, for a 768×1024 image. The proposed technique is somewhat more complex than the original, but not dramatically so. Most of the gap is due to the costly nonlocal denoising. The decision phase, instead, takes much less than expected. Thanks to the efficient optimization method, it requires less than 12 sec. without subsampling, which becomes just 0.31 sec. with subsampling, less than 5% of the overall CPU time. After the final upsampling, some inexpensive morphological filtering is used to slightly enlarge and smooth all map contours.

Table 7.1 *CPU time (sec.) for the original and proposed techniques.*

	CFAR-Mihcak		Bayes-BM3D	
	mean	st.dev.	mean	st.dev.
Denoising	0.45	0.03	5.45	0.12
Index Field Computation	1.30	0.01	1.30	0.01
Decision/Optimization	0.01	0.00	0.31	0.14
Morphological Filtering	0.03	0.00	0.04	0.00
Overall	1.79	0.03	7.10	0.21

7.4.6 Comparison with other reference techniques

To complete our analysis, we show here the results of an experiment designed to compare performance with some state-of-the-art techniques, namely A-DJPG [17] (based on JPEG artifacts), and CFA-Loc [97] (based on the analysis of the CFA), for which the codes are publicly available. In order to fairly compare techniques based on very different principles and hypothesis, a pretty realistic setting was designed. A test set of 300 images of size 768×1024, 75 for each of our test cameras, were created with realistic forgeries, some of them drawn from the web, and others designed by ourselves. We included an equal number of small, medium, and large forgeries, classified based on the largest circle included in the forgery area, with radius $r \leq 64$, $r \in [65, 128]$ and $r \in [129, 192]$, respectively. Both uncompressed and compressed ($Q > 75$) forgeries were used, half of them resampled, the other half either rescaled or rotated.

Figure 7.8 shows ROCs obtained at pixel level on the test set. To evaluate robustness of the algorithms, we considered also the case in which all tampered images were JPEG compressed with quality factor 90. In both cases, Bayes-BM3D performs better than CFAR-Mihcak, which in turn performs slightly better than A-DJPG, and much better than CFA-Loc. In A-DJPG, we used all 64 DCT coefficients to generate the likelihood map, since 6 coefficients were not enough to obtain a good performance [17]. It is also interesting to note that while all techniques exhibit a worse performance on compressed images, the impairment is much stronger for A-DJPG and CFA-Loc than for the PRNU-based techniques.

Although a thorough comparative analysis would be certainly interesting, it is out of the scope of this work. It is worth pointing out that the proposed technique improves upon the well-known and extensively tested technique in [41], and can be therefore expected to inherit its good and robust performance.

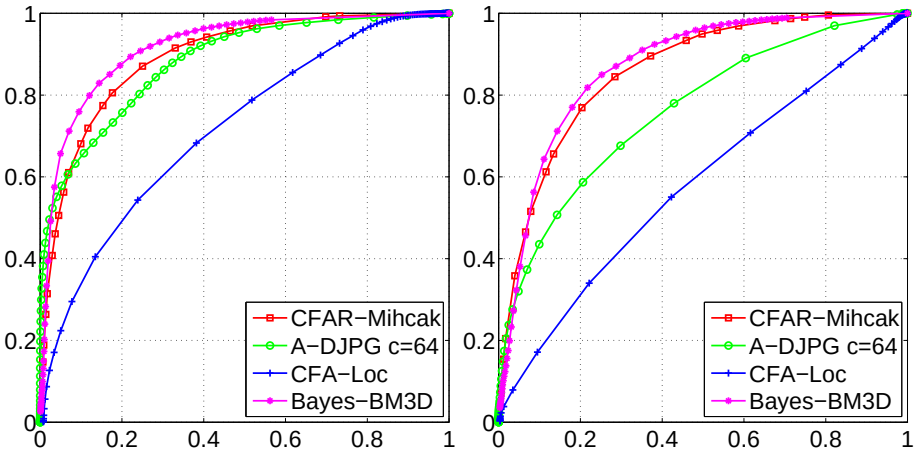

(a) Uncompressed images. (b) Compressed images.

Figure 7.8 *Pixel-level ROCs for PRNU-based and reference techniques.*

7.5 Conclusions

Image forgery detection is becoming more challenging by the day, due to the unrelenting advances in image processing. PRNU-based forgery detection is one of the most promising approaches for this task. Provided that one has the opportunity to estimate the camera PRNU, all kinds of forgeries can be dealt with in a uniform manner and with a consistent good performance. Here, we improve upon the seminal PRNU-based forgery detection technique proposed in [41] by recasting the problem in a Bayesian framework, and by modeling the decision variables as a Markov random field, thus accounting for their spatial dependencies. In addition, we resort to state-of-the-art signal and image processing tools: nonlocal denoising to improve estimation of noise residual, and convex optimization to reach a globally optimal solution in a limited time. As a result, the proposed technique provides a significant and consistent performance gain over the original one, especially in terms of object-level detection ability, the main parameter of interest for the applications. A modified version of the proposed algorithm together with some other forgery detection tools [77], allowed the GRIP team to win the First IEEE IFS-TC Image Forensics Challenge.

You may not end up where you thought you'd be, but you'll end up right where you're meant to be.

ANONYMOUS

CONCLUSION AND PERSPECTIVES

In this thesis, we have designed a number of convex optimization algorithms to address three problems in image restoration, machine learning, and digital forensics. The common thread among these algorithms is the use of a novel epigraphical splitting technique for dealing with nonlinear convex constraints. Experiments carried out for each application have illustrated the efficiency and the performance of the proposed approaches w.r.t. state-of-the-art solutions.

CONCLUSION

In the first part of the thesis, we have proposed a new epigraphical splitting technique to deal with a class of nonlinear convex constraints. The proposed technique allows us to reduce the complexity of optimization algorithms when a constraint involves the sum of absolute values raised to a power $q \geq 1$, distance functions to a convex set, Euclidean norms, infinity norms, and max functions. To demonstrate the efficiency of our approach, we have focused on mixed-norm constraints that enforce a sparse regularization in image recovery problems. In this context, the obtained results indicate that algorithms based on the epigraphical splitting are faster (in terms of execution time) than equivalent algorithms based on the direct computation of the projections via standard iterative solutions. A summary of these findings was published in [51].

Although the proposed splitting turns out to be very efficient when the epigraphical projections can be quickly evaluated, the same efficiency cannot be claimed when there is no available expression for such projections. To facilitate the use of the epigraphical splitting, in the first part of the thesis, we have also turned our attention to outer-approximated constraints based on piecewise-affine functions. In this context, we have tackled constraints involving the Kullback-Leibler divergence and the logistic loss. The obtained results indicate both the good performance of the proposed approximation and the efficiency of the epigraphical splitting. These findings were partially published in [50].

In the second part of the thesis, grounded on the aforementioned epigraphical splitting, we have brought three contributions in the context of multicomponent image recovery, sparse multiclass SVM learning, and PRNU-based image forgery detection, as summarized in the following.

(i). We have proposed a new regularization for multicomponent images that combines nonlocal total variation and structure tensor. The obtained results demonstrate the better performance of structure tensor and nonlocal

gradients. They also indicate that the epigraphical splitting leads to faster algorithms (in terms of the execution time) than those based on the direct computation of the projections via standard iterative solutions. This contribution was published in [53], and the toolbox is publicly available.[1]

(ii). We have proposed two efficient algorithms for learning a sparse multiclass SVM. Our approach makes it possible to minimize a criterion involving the multiclass hinge loss and a sparsity-inducing regularization. In the literature, such a criterion is typically approximated by replacing the hinge loss with a smooth penalty, such as the quadratic hinge loss or the logistic loss. The obtained results show that the use of the hinge loss, rather than an approximation, leads to a slightly better classification accuracy, while the proposed algorithms are often faster than state-of-the-art solutions, with the one based on the epigraphical splitting being the most efficient. This contribution is, at the moment, under review [52].

(iii). We have proposed a new approach for PRNU-based image forgery detection by recasting the problem in a Bayesian framework, and by modeling the decision variables as a Markov random field, thus accounting for their spatial dependencies. In addition, we have resorted to nonlocal denoising to improve estimation of noise residual, and proximal methods based on the epigraphical splitting to reach a globally optimal solution in a limited time. As a result, the proposed technique provides a significant and consistent performance gain over the original one, especially in terms of object-level detection ability, the main parameter of interest for the applications. This contribution was published in [49].

Perspectives

Despite the present advances, the work in this thesis opens up a number of interesting perspectives from an applicative standpoint, as discussed next.

- The regularization based on nonlocal structure tensor leads to competitive results in the context of spectral imaging, and thus it would be interesting to consider other applications, such as the recovery of video signals or volumetric images. Another possibility could be the investigation of techniques for reducing the high computational burden due to the huge number of spectral bands. This goal may be achieved, for example, by processing only the most informative bands selected by means of PCA, as already done in [240] for the super-resolution of hyperspectral images. Other solutions may involve the approximation of the singular value decomposition, which also take a significant amount of time. Last but not least, ad-hoc implementations on massively parallel architectures should allow one to further accelerate the proposed algorithms [102].

- We believe that sparse multiclass SVM learning is one of the most promising topics in the thesis. To the best of our knowledge, there exists no real alternative for dealing with the exact formulation of the multiclass hinge

[1] `perso.ens-lyon.fr/nelly.pustelnik/Software/Toolbox_MultiResto_v-1.0.zip`

loss, as the latter is typically approximated with a smooth penalty, such as the quadratic hinge loss or the logistic loss. Currently, one of the main limitations of sparse learning is the lack of methods for dealing with a large amount of training data, or a huge number of classes. Standard SVMs circumvent this difficulty by resorting to Lagrangian duality techniques [78], leading to a constrained quadratic formulation that can be efficiently decomposed into smaller problems [19, 183], or even approximated through cutting plane approaches, in order to address scenarios with a lot of training data and thousands (or even an infinite number) of classes [126, 213]. However, the same reasoning cannot be applied to sparse SVMs, as the dual approach brings no advantages whatsoever. One way to overcome this difficulty consists of resorting to block-coordinate descent methods [20], but they impose restrictive assumptions on the problem to be solved (this is the main reason why quadratic approximations of the hinge loss are used). Another approach is given by stochastic methods [242].

- There is still much room for improvements in the context of PRNU-based image forgery detection. A notable example is the design of a better and more robust predictor, in order to reduce the False Alarm rate. Our major goal for future research, however, is to improve spatial resolution, allowing for the detection of smaller forgeries. Prior work on this topic [48] showed that image-level segmentation can help increasing resolution in some suitable cases, but segmentation itself is a very challenging and unreliable process. We are currently working towards a new version of this algorithm based on soft segmentation [46].

- We are currently working on other applications that can benefit from the epigraphical splitting. One of the most interesting concerns the optimal rate allocation in predictive video coding, which is challenging because of the dependencies between frames induced by motion compensation. In a preliminary work [98], we derived an analytical rate-distortion model that explicitly takes into account the dependencies between frames, allowing us to formulate the optimal rate allocation as a convex optimization problem that we solved (exactly and efficiently) via proximal methods. An interesting extension of this work consists of considering a constrained formulation of the rate-allocation problem, so that the user can have a direct control on the distortion of coded frames, the main parameter of interest for this application. The epigraphical splitting thus arises as a natural solution for this kind of problems, and we expect to obtain a performance in line with the results presented in this thesis.

In addition to these avenues of research, we feel that the approximation method presented in Chapter 4 has not been developed yet to its full potential, and it would be interesting to investigate the connections with bundle methods [57, 73].

Bibliography

[1] M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo. "Fast image recovery using variable splitting and constrained optimization". In: *IEEE Trans. Image Process.* 19.9 (2010), pages 2345–2356 (cited on pages 6, 16, 62).

[2] M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo. "An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems". In: *IEEE Trans. Image Process.* 20.3 (Mar. 2011), pages 681–695 (cited on pages 4, 5, 62).

[3] M. Aharon, M. Elad, and A. Bruckstein. "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation". In: *IEEE Trans. Signal Process.* 54.11 (Nov. 2006), pages 4311–4322 (cited on page 27).

[4] M. Aizerman, E. Braverman, and L. Rozonoer. "Theoretical foundations of the potential function method in pattern recognition learning". In: *Autom. Remote Control* 25 (1964), pages 821–837 (cited on page 71).

[5] I. Amerini, L. Ballan, R. Caldelli, A. D. Bimbo, and G. Serra. "A SIFT-Based Forensic Method for Copy and Move Attack Detection and Transformation Recovery". In: *IEEE Trans. Inf. Forensics Security* 6.3 (2011), pages 1099–1110 (cited on page 87).

[6] H. Attouch, G. Buttazzo, and G. Michaille. *Variational Analysis in Sobolev and BV Spaces: Applications to PDEs and Optimization.* MOS-SIAM Series on Optimization, 2006 (cited on pages xxxii, 53, 63, 64).

[7] J.-F. Aujol. "Some first-order algorithms for total variation based image restoration". In: *J. Math. Imag. Vis.* 34.3 (July 2009), pages 307–327 (cited on page 58).

[8] J.-F. Aujol, G. Gilboa, T. Chan, and S. Osher. "Structure-Texture Image Decomposition - Modeling, Algorithms, and Parameter Selection". In: *Int. J. Comp. Vis.* 67.1 (Apr. 2006), pages 111–136 (cited on page 4).

[9] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. "Optimization with Sparsity-Inducing Penalties". In: *Foundations and Trends in Machine Learning* 4.1 (2012), pages 1–106 (cited on pages 47, 71, 75).

[10] S. Battiato, G. M. Farinella, G. M. Messina, and G. Puglisi. "Robust Image Alignment for Tampering Detection". In: *IEEE Trans. Inf. Forensics Security* 7.4 (2012), pages 1105–1117 (cited on page 87).

[11] F. Baus, M. Nikolova, and G. Steidl. "Fully smoothed L1-TV models: Bounds for the minimizers and parameter choice". In: *J. Math. Imaging Vis.* 48.2 (Feb. 2013), pages 295–307 (cited on page 5).

[12] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, Jan. 2011 (cited on pages 6–12, 17, 76).

[13] S. Bayram, H. T. Sencar, and N. Memon. "Efficient Sensor Fingerprint Matching Through Fingerprint Binarization". In: *IEEE Trans. Inf. Forensics Security* 7.4 (Aug. 2012), pages 1404–1413 (cited on page 89).

[14] A. Ben-Tal and M. Teboulle. "A smoothing technique for nondifferentiable optimization problems". In: *Lecture Notes in Mathematics* 1405 (1989), pages 1–11 (cited on page 5).

[15] K. Bernard, Y. Tarabalka, J. Angulo, J. Chanussot, and J. A. Benediktsson. "Spectral-Spatial Classification of Hyperspectral Data Based on a Stochastic Minimum Spanning Forest Approach". In: *IEEE Trans. Image Process.* 21.4 (Apr. 2012), pages 2008–2021 (cited on page 55).

[16] J. Besag. "On the statistical analysis of dirty pictures". In: *J. R. Stat. Soc. B* 48 (1986), pages 259–302 (cited on pages 94, 96).

[17] T. Bianchi and A. Piva. "Image Forgery Localization via Block-Grained Analysis of JPEG Artifacts". In: *IEEE Trans. Inf. Forensics Security* 7 (June 2012), pages 1003–1017 (cited on pages 87, 109).

[18] P. Blomgren and T. Chan. "Color TV: total variation methods for restoration of vector-valued images". In: *IEEE Trans. Image Process.* 7.3 (Mar. 1998), pages 304–309 (cited on page 53).

[19] M. Blondel, A. Fujino, and N. Ueda. "Large-scale Multiclass Support Vector Machine Training via Euclidean Projection onto the Simplex". In: *Proc. of ICPR*. Stockholm, Sweden, Aug. 2014, pages 1289–1294 (cited on pages 71, 76, 113).

[20] M. Blondel, K. Seki, and K. Uehara. "Block Coordinate Descent Algorithms for Large-scale Sparse Multiclass Classification". In: *Mach. Learn.* 93.1 (Oct. 2013), pages 31–52 (cited on pages 72, 78, 79, 113).

[21] P. Bouboulis, K. Slavakis, and S. Theodoridis. "Adaptive learning in complex reproducing kernel Hilbert spaces employing Wirtinger's subgradients". In: *IEEE Trans. Neural Netw.* 23.3 (Mar. 2012), pages 425–438 (cited on page 11).

[22] A. C. Bovik. *Handbook of Image and Video Processing*. Series in Communications, Networking, and Multimedia. Academic Press, 2005 (cited on page 26).

[23] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers". In: *Foundations and Trends in Machine Learning* 8.1 (2011), pages 1–122 (cited on pages 16, 47).

[24] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambrige University Press, 2004 (cited on pages 48, 74, 75).

[25] Y. Boykov, O. Veksler, and R. Zabih. "Fast approximate energy minimization via graph cuts". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 23 (2001), pages 1222–1239 (cited on page 96).

[26] P. S. Bradley and O. L. Mangasarian. "Feature Selection via Concave Minimization and Support Vector Machines". In: *Proc. of ICML*. Madison, USA, July 1998, pages 82–90 (cited on page 72).

[27] K. Bredies, K. Kunisch, and T. Pock. "Total Generalized Variation". In: *SIAM J. Imaging Sci.* 3.3 (Sept. 2010), pages 492–526 (cited on page 27).

[28] X. Bresson and T. F. Chan. "Fast dual minimization of the vectorial total variation norm and applications to color image processing". In: *Inverse Probl. Imag.* 2.4 (Nov. 2008), pages 455–484 (cited on pages xxxii, 27, 53, 56, 58, 63, 64).

[29] L. M. Briceño-Arias and P. L. Combettes. "A monotone + skew splitting model for composite monotone inclusions in duality". In: *SIAM J. Opt.* 21.4 (Oct. 2011), pages 1230–1250 (cited on page 6).

[30] L. M. Briceño-Arias, P. L. Combettes, J.-C. Pesquet, and N. Pustelnik. "Proximal algorithms for multicomponent image recovery problems". In: *J. Math. Imag. Vis.* 41.1 (Sept. 2011), pages 3–22 (cited on pages 16, 56, 62).

[31] J. Bruna and S. Mallat. "Invariant Scattering Convolution Networks". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 35.8 (Aug. 2013), pages 1872–1886 (cited on page 78).

[32] A. Buades, B. Coll, and J. Morel. "A review of image denoising algorithms, with a new one". In: *Multiscale Model. and Simul.* 4.2 (2005), pages 490–530 (cited on pages 27, 58, 98).

[33] E. J. Candès, X Li, Y. Ma, and J. Wright. "Robust principal component analysis?" In: *J. ACM* 58.3 (May 2011), 11:1–11:37 (cited on page 56).

[34] E. J. Candès, M. B. Wakin, and S. Boyd. "Enhancing Sparsity by Reweighted ℓ_1 Minimization". In: *J. Fourier Anal. Appl.* 14.5 (Dec. 2008), pages 877–905 (cited on pages 4, 15, 72).

[35] Y. Censor, W. Chen, P. L. Combettes, R. Davidi, and G. T. Herman. "On the effectiveness of projection methods for convex feasibility problems with linear inequality constraints". In: *Comput. Optim. Appl.* 51.3 (2012), pages 1065–1088 (cited on page 5).

[36] A. Chambolle, D. Cremers, and T. Pock. "A Convex Approach to Minimal Partitions". In: *SIAM J. Imaging Sci.* 5.4 (2012), pages 1113–1158 (cited on page 4).

[37] A. Chambolle and T. Pock. "A first-order primal-dual algorithm for convex problems with applications to imaging". In: *J. Math. Imag. Vis.* 40.1 (May 2011), pages 120–145 (cited on page 6).

[38] T. F. Chan, S. Esedoglu, and M. Nikolova. "Algorithms for finding global minimizers of image segmentation and denoising models". In: *SIAM J. App. Math.* 66.5 (2006), pages 1632–1648 (cited on page 96).

[39] C. Chaux, A. Benazza-Benyahia, J.-C. Pesquet, and Duval. "Wavelet transform for the denoising of multivariate images". In: *Multivariate Image Processing*. Edited by C. Collet, J. Chanussot, and K. Chehdi. John Wiley & Sons Inc., 2010, pages 203–237 (cited on page 27).

[40] C. Chaux, P. L. Combettes, J.-C. Pesquet, and V. R. Wajs. "A variational formulation for frame-based inverse problems". In: *Inverse Probl.* 23.4 (June 2007), pages 1495–1518 (cited on pages 4, 6, 9, 18).

[41] M. Chen, J. Fridrich, M. Goljan, and J. Lukas. "Determining Image Origin and Integrity Using Sensor Noise". In: *IEEE Trans. Inf. Forensics Security* 3.1 (Mar. 2008), pages 74–90 (cited on pages xxxviii, 88–93, 96, 109, 110).

[42] Y. Chen and A. O. Hero. "Recursive $\ell_{1,\infty}$ lasso". In: *IEEE Trans. Signal Process.* 60.8 (Aug. 2012), pages 3978–3987 (cited on page 16).

[43] Y. Chen and V. L. L. Thing. "A Study on the Photo Response Non-Uniformity Noise Pattern based Image Forensics in Real-World Applications". In: *WORLDCOMP*. 2012 (cited on page 88).

[44] Y.-L. Chen and C.-T. Hsu. "Detecting Recompression of JPEG Images via Periodicity Analysis of Compression Artifacts for Tampering Detection". In: *IEEE Trans. Inf. Forensics Security* 6.2 (June 2011), pages 396–406 (cited on page 87).

[45] Q. Cheng, H. Shen, L. Zhang, and P. Li. "Inpainting for Remotely Sensed Images With a Multichannel Nonlocal Total Variation Model". In: *IEEE Trans. Geosci. Remote Sens.* 52.1 (Jan. 2014), pages 175–187 (cited on pages xxxiii, 56, 58, 63, 65, 66).

[46] G. Chierchia, D. Cozzolino, G. Poggi, C. Sansone, and L. Verdoliva. "Guided filtering for PRNU-based localization of small-size image forgeries". In: *Proc. of ICASSP*. 2014 (cited on page 113).

[47] G. Chierchia, S. Parrilli, G. Poggi, C. Sansone, and L. Verdoliva. "On the influence of denoising in PRNU based forgery detection". In: *Proc. of ACM MiFor*. Florence, Italy, 2010, pages 117–122 (cited on pages 88, 93).

[48] G. Chierchia, S. Parrilli, G. Poggi, L. Verdoliva, and C. Sansone. "PRNU-based detection of small-size image forgeries". In: *Int. Conf. Digital Signal Process.* 2011, pages 1–6 (cited on pages 89, 113).

[49] G. Chierchia, G. Poggi, C. Sansone, and L. Verdoliva. "A Bayesian-MRF Approach for PRNU-based Image Forgery Detection". In: *IEEE Trans. Inf. Forensics Security* 9.4 (Apr. 2014), pages 554–567 (cited on pages xii, 112).

[50] G. Chierchia, N. Pustelnik, J.-C. Pesquet, and B. Pesquet-Popescu. "A proximal approach approach for constrained cosparse modelling". In: *Proc. of ICASSP*. Kyoto, Japan, Mar. 2012 (cited on pages xi, 111).

[51] G. Chierchia, N. Pustelnik, J.-C. Pesquet, and B. Pesquet-Popescu. "Epigraphical Projection and Proximal Tools for Solving Constrained Convex Optimization Problems". In: *Signal, Image and Video Processing* (July 2014) (cited on pages xi, xxxii, 63, 64, 111).

[52] G. Chierchia, N. Pustelnik, J.-C. Pesquet, and B. Pesquet-Popescu. "A Proximal Approach for Sparse Multiclass SVM". In: *arxiv.org/abs/1501.03669* (Feb. 2015). Preprint (cited on pages xii, 112).

[53] G. Chierchia, N. Pustelnik, B. Pesquet-Popescu, and J.-C. Pesquet. "A Nonlocal Structure Tensor Based Approach for Multicomponent Image Recovery Problems". In: *IEEE Trans. Image Process.* 23.12 (Dec. 2014), pages 5531–5544 (cited on pages xii, 112).

[54] E. Chouzenoux, J. Idier, and S. Moussaoui. "A Majorize-Minimize Strategy for Subspace Optimization Applied to Image Restoration". In: *IEEE Trans. Image Process.* 20.6 (June 2011), pages 1517–1528 (cited on page 5).

[55] E. Chouzenoux, M. Legendre, S. Moussaoui, and J. Idier. "Fast Constrained Least Squares Spectral Unmixing using Primal-Dual Interior Point Optimization". In: *IEEE J. Sel. Topics Appl. Earth Observ. in Remote Sens.* (2013), pages 59–69 (cited on page 54).

[56] V. Christlein, C. Riess, J. Jordan, C. Riess, and E. Angelopoulou. "An Evaluation of Popular Copy-Move Forgery Detection Approaches". In: *IEEE Trans. Inf. Forensics Security* 7.6 (Dec. 2012), pages 1841–1854 (cited on page 87).

[57] P. L. Combettes. "Strong convergence of block-iterative outer approximation methods for convex optimization". In: *SIAM J. Control Optim.* 38.2 (Feb. 2000), pages 538–565 (cited on pages 40, 113).

[58] P. L. Combettes. "A block-iterative surrogate constraint splitting method for quadratic signal recovery". In: *IEEE Trans. Signal Process.* 51.7 (July 2003), pages 1771–1782 (cited on page 11).

[59] P. L. Combettes, L. Condat, J.-C. Pesquet, and B. C. Vũ. "A forward-backward view of some primal-dual optimization methods in image recovery". In: *Proc. of ICIP*. Paris, France, Oct. 2014 (cited on pages 6, 8).

[60] P. L. Combettes, D. Dũng, and B. Công Vũ. "Dualization of Signal Recovery Problems". In: *Set-Valued Var. Anal.* 18 (Dec. 2010), pages 373–404 (cited on page 7).

[61] P. L. Combettes and J.-C. Pesquet. "Image restoration subject to a total variation constraint". In: *IEEE Trans. Image Process.* 13.9 (Sept. 2004), pages 1213–1222 (cited on pages 5, 13, 26, 55, 67, 97).

[62] P. L. Combettes and J.-C. Pesquet. "A Douglas-Rachford Splitting Approach to Nonsmooth Convex Variational Signal Recovery". In: *IEEE J. Sel. Topics Signal Process.* 1.4 (Dec. 2007), pages 564–574 (cited on page 21).

[63] P. L. Combettes and J.-C. Pesquet. "Proximal thresholding algorithm for minimization over orthonormal bases". In: *SIAM J. Opt.* 18.4 (Nov. 2007), pages 1351–1376 (cited on page 18).

[64] P. L. Combettes and J.-C. Pesquet. "A proximal decomposition method for solving convex variational inverse problems". In: *Inverse Probl.* 24.6 (Dec. 2008) (cited on pages 15, 19, 62).

[65] P. L. Combettes and J.-C. Pesquet. "Proximal splitting methods in signal processing". In: *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Edited by H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, et al. New York: Springer-Verlag, 2011, pages 185–212 (cited on pages 6, 9, 28, 30, 62, 76).

[66] P. L. Combettes and J.-C. Pesquet. "Primal-dual splitting algorithm for solving inclusions with mixtures of composite, Lipschitzian, and parallel-sum type monotone operators". In: *Set-Valued Var. Anal.* 20.2 (June 2012), pages 307–330 (cited on pages xl, 6, 8, 28, 29, 97).

[67] P. L. Combettes and H. J. Trussell. "The use of noise properties in set theoretic estimation". In: *IEEE Trans. Signal Process.* 39.7 (July 1991), pages 1630–1641 (cited on page 5).

[68] P. L. Combettes and V. R. Wajs. "Signal Recovery by proximal forward-backward splitting". In: *Multiscale Model. and Simul.* 4.4 (Nov. 2005), pages 1168–1200 (cited on pages 6, 49).

[69] L. Condat. "A primal-dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms". In: *J. Optim. Theory Appl.* 158.2 (2013), pages 460–479 (cited on pages 6, 8, 45, 49, 72, 76).

[70] L. Condat. "A Generic Proximal Algorithm for Convex Optimization - Application to Total Variation Minimization". In: *IEEE Signal Proc. Letters* 21.8 (Aug. 2014), pages 1054–1057 (cited on page 63).

[71] L. Condat. "Fast Projection onto the Simplex and the l1 Ball". In: (2014). http://hal.archives-ouvertes.fr/hal-01056171 (cited on pages 10, 12, 13, 28, 59, 62, 67, 76, 97).

[72] L. Condat and S. Mosaddegh. "Joint Demosaicking and Denoising by Total Variation Minimization". In: *Proc. of ICIP*. Orlando, USA, Sept. 2012 (cited on page 63).

[73] R. Correa and C. Lemaréchal. "Convergence of some algorithms for convex minimization". In: *Math. Program.* 62.1–3 (1993), pages 261–275 (cited on page 113).

[74] L. Corrias. "Fast Legendre-Fenchel Transform and Applications to Hamilton-Jacobi Equations andConservation Laws". In: *SIAM J. Numer. Anal.* 33.4 (Aug. 1996), pages 1534–1558 (cited on page 42).

[75] C. Cortes and V. Vapnik. "Support-Vector Networks". In: *Mach. Learn.* 20.3 (Sept. 1995), pages 273–297 (cited on pages 71, 72).

[76] T. M. Cover. "Geometrical and Statistical Properties of Systems of Linear Inequalities with Applications in Pattern Recognition". In: *IEEE Trans. on Electronic Computers* EC-14.3 (June 1965), pages 326–334 (cited on page 73).

[77] D. Cozzolino, D. Gragnaniello, and L. Verdoliva. *A novel framework for image forgery localization.* Technical report. IEEE IFS-TC Image Forensics Challenge - (Phase 2). Nov. 2013. URL: `http://arxiv.org/pdf/1311.6932v1.pdf` (cited on page 110).

[78] K. Crammer and Y. Singer. "On the Algorithmic Implementation of Multiclass Kernel-based Vector Machines". In: *J. Mach. Learn. Res.* 2 (2001), pages 265–392 (cited on pages 13, 16, 71, 72, 74, 76, 113).

[79] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. "Image denoising by sparse 3D transform-domain collaborative filtering". In: *IEEE Trans. Image Process.* 16.8 (Aug. 2007), pages 2080–2095 (cited on pages 27, 98).

[80] A. Danielyan, V. Katkovnik, and K. Egiazarian. "BM3D Frames and Variational Image Deblurring". In: *IEEE Trans. Image Process.* 21.4 (Apr. 2012), pages 1715–1728 (cited on pages 4, 15, 27).

[81] I. Daubechies, M. Defrise, and C. De Mol. "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint". In: *Comm. Pure Applied Math.* 57.11 (Nov. 2004), pages 1413–1457 (cited on page 4).

[82] C. Deledalle, S. Vaiter, G. Peyré, and J. Fadili. "Stein Unbiased GrAdient estimator of the Risk (SUGAR) for multiple parameter selection". In: *SIAM J. Imaging Sci.* 7.4 (2014), pages 2448–2487 (cited on page 5).

[83] C. D'Elia, G. Poggi, and G. Scarpa. "A tree-structured Markov random field model for Bayesian image segmentation". In: *IEEE Trans. Image Process.* 12.10 (Oct. 2003), pages 1259–1273 (cited on page 94).

[84] S. Di Zenzo. "A Note on the Gradient of a Multi-Image". In: *Comput. Vision Graph. Image Process.* 33.1 (Jan. 1986), pages 116–125 (cited on page 53).

[85] C. Ding, D. F. Sun, and K.-C. Toh. "An introduction to a class of matrix cone programming". In: *Math. Program.* (Dec. 2012), pages 1–39 (cited on page 20).

[86] N. Dobigeon, S. Moussaoui, M. Coulon, J.-Y. Tourneret, and A. O. Hero. "Joint Bayesian endmember extraction and linear unmixing for hyperspectral imagery". In: *IEEE Trans. Signal Process.* 57.11 (Nov. 2009), pages 4355–4368 (cited on page 54).

[87] D. L. Donoho and I. M. Johnstone. "Adapting to unknown smoothness via wavelet shrinkage". In: *J. Am. Statist. Assoc.* 90.432 (1995), pages 1200–1224 (cited on page 5).

[88] J. Duchi and Y. Singer. "Boosting with Structural Sparsity". In: *Proc. of ICML*. Montreal, Canada, June 2009, pages 297–304 (cited on pages 71, 72).

[89] V. Duval, J.-F. Aujol, and L. Vese. "Projected Gradient Based Color Image Decomposition". In: *Proc. of SSVM*. Volume 5567. Springer Verlag, 2009, pages 295–306 (cited on page 53).

[90] J. Eckstein. "Parallel alternating direction multiplier decomposition of convex programs". In: *J. Optim. Theory Appl.* 80.1 (Jan. 1994), pages 39–62 (cited on page 62).

[91] J. Eckstein and D. P. Bertsekas. "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators". In: *Math. Program.* 55 (3 1992), pages 293–318 (cited on pages 16, 62).

[92] M. Elad and M. Aharon. "Image denoising via sparse and redundant representations over learned dictionaries". In: *IEEE Trans. Image Process.* 15.12 (Dec. 2006), pages 3736–3745 (cited on page 27).

[93] Y. C. Eldar. "Generalized SURE for Exponential Families: Applications to Regularization". In: *IEEE Trans. Signal Process.* 57.2 (Feb. 2009), pages 471–481 (cited on page 5).

[94] G. Ely, S. Aeron, and E. L. Miller. "Exploiting Structural Complexity for Robust and Rapid Hyperspectral Imaging". In: *Proc. of ICASSP*. Vancouver, Canada, June 2013 (cited on page 56).

[95] J. M. Fadili and G. Peyré. "Total Variation Projection With First Order Schemes". In: *IEEE Trans. Image Process.* 20.3 (Mar. 2011), pages 657–669 (cited on pages 59, 62).

[96] M. J. Fadili, J.-L. Starck, and F. Murtagh. "Inpainting and Zooming Using Sparse Representations". In: *Comput. J.* 52.1 (Jan. 2009), pages 64–79 (cited on page 53).

[97] P. Ferrara, T. Bianchi, A. D. Rosa, and A. Piva. "Image Forgery Localization via Fine-Grained Analysis of CFA Artifacts". In: *IEEE Trans. Inf. Forensics Security* 7.5 (2012), pages 1566–1577 (cited on pages 88, 109).

[98] N. Fiengo, G. Chierchia, M. Cagnazzo, and B. Pesquet-Popescu. "A convex-optimization framework for frame-level optimal rate allocation in predictive video coding". In: *Proc. of ICASSP*. Florence, Italy, 2014 (cited on pages xli, 113).

[99] A. Foi and G. Boracchi. "Foveated self-similarity in nonlocal image filtering". In: *Proc. of SPIE*. Volume 8291. Burlingame (CA), USA, Jan. 2012 (cited on pages 27, 58).

[100] J. Fridrich. "Sensor Defects in Digital Image Forensic". In: *Digital Image Forensics* (2012), pages 179–218 (cited on pages 88, 89).

[101] H. Fu and X. Cao. "Forgery Authentication in Extreme Wide-Angle Lens Using Distortion Cue and Fake Saliency Map". In: *IEEE Trans. Inf. Forensics Security* 7.4 (Aug. 2012), pages 1301–1314 (cited on page 88).

[102] R. Gaetano, G. Chierchia, and B. Pesquet-Popescu. "Parallel Implementations of a Disparity Estimation Algorithm based on a Proximal Splitting Method". In: *Proc. of VCIP*. San Diego, USA, 2012 (cited on pages 14, 37, 112).

[103] S. Geman and D. Geman. "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 6.6 (Nov. 1984), pages 721–741 (cited on page 94).

[104] G. Gilboa and S. Osher. "Nonlocal Linear Image Regularization and Supervised Segmentation". In: *Multiscale Model. and Simul.* 6.2 (2007), pages 595–630 (cited on pages 27, 58).

[105] G. Gilboa and S. Osher. "Nonlocal Operators with Applications to Image Processing". In: *Multiscale Model. and Simul.* 7.3 (2009), page 1005 (cited on pages 25, 27, 58).

[106] T. Gloe, S. Pfennig, and M. Kirchner. "Unexpected Artefacts in PRNU-Based Camera Identification : A Dresden Image Database Case-Study". In: *Proc. of ACM MM&Sec.* 2012, pages 109–114 (cited on page 89).

[107] M. Golbabaee and P. Vandergheynst. "Hyperspectral Image Compressed Sensing via Low-Rank and Joint-Sparse Matrix Recovery". In: *Proc. of ICASSP*. Kyoto, Japan, Mar. 2012 (cited on pages 53, 54, 56).

[108] B. Goldluecke, E. Strekalovskiy, and D. Cremers. "The Natural Vectorial Total Variation Which Arises from Geometric Measure Theory". In: *SIAM J. Imaging Sci.* 5.2 (2012), pages 537–563 (cited on pages xxxii, 53, 56, 58, 63, 64).

[109] M. Goljan and J. Fridrich. "Digital camera identification from images – Estimating false acceptance probability". In: *Proc. of IWDW*. 2008 (cited on page 89).

[110] M. Goljan and J. Fridrich. "Sensor-fingerprint based identification of images corrected for lens distortion". In: *Proc. of SPIE*. vol. 8303. 2012 (cited on page 89).

[111] M. Goljan, J. Fridrich, and T. Filler. "Managing a large database of camera fingerprints". In: *Proc. of SPIE*. Volume 7541. 2010 (cited on page 89).

[112] T. R. Golub, D. K. Slonim, P. Tamayo, C. Huard, M. Gaasenbeek, et al. "Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring". In: *Science* 286.5439 (1999), pages 531–537 (cited on page 78).

[113] S. Harizanov, J.-C. Pesquet, and G. Steidl. "Epigraphical Projection for Solving Least Squares Anscombe Transformed Constrained Optimization Problems". In: *Proc. of SSVM*. Volume 7893. 2013, pages 125–136 (cited on page 13).

[114] G. E. Healey and R. Kondepudy. "Radiometric CCD camera calibration and noise estimation". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 16.3 (Mar. 1994), pages 267–276 (cited on pages 88, 90).

[115] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms, Part I: Fundamentals*. Volume 305. Springer-Verlag, 1996 (cited on page 5).

[116] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex analysis and minimization algorithms, Part II : Advanced theory and bundle methods*. Volume 305. Springer-Verlag, 2001 (cited on page 40).

[117] D. Hsu, S. M. Kakade, and T. Zhang. "Robust matrix decomposition with sparse corruptions". In: *IEEE Trans. Inf. Theory* 57.11 (Nov. 2011), pages 7221–7234 (cited on page 56).

[118] Y. Hu and M. Jacob. "Higher Degree Total Variation (HDTV) Regularization for Image Recovery". In: *IEEE Trans. Image Process.* 21.5 (May 2012), pages 2559–2571 (cited on page 27).

[119] Y. Hu and M. Jacob. "Generalized Higher Degree Total Variation (HDTV) Regularization". In: *IEEE Trans. Image Process.* 23.6 (June 2014), pages 2423–2435 (cited on page 27).

[120] Y. Hu, C. Li, Z. Lai, and S. Zhang. "Fast Camera Fingerprint Search Algorithm for Source Camera". In: *Int. Symp. on Comm. Control Signal Process.* 2012, pages 2–4 (cited on page 89).

[121] F. J. Huang and Y. LeCun. "Large-scale Learning with SVM and Convolutional for Generic Object Categorization". In: *Proc. of CVPR*. New York, USA, June 2006, pages 284–291 (cited on page 71).

[122] Y. Huang, L. Moisan, M. Ng, and T. Zeng. "Multiplicative noise removal via a learned dictionary". In: *IEEE Trans. Image Process.* 21.11 (Nov. 2012), pages 4534–4543 (cited on page 27).

[123] P. J. Huber. "Robust Estimation of a Location Parameter". In: *Ann. Math. Statist.* 35.1 (1964), pages 73–101 (cited on page 5).

[124] L. Jacques, L. Duval, C. Chaux, and G. Peyré. "A Panorama on Multiscale Geometric Representations, Intertwining Spatial, Directional and Frequency Selectivity". In: *Signal Process.* 91.12 (Dec. 2011), pages 2699–2730 (cited on pages 25, 27).

[125] T. Joachims. "Text Categorization with Suport Vector Machines: Learning with Many Relevant Features". In: *Proc. of ECML*. Chemnitz, Germany, Apr. 1998, pages 137–142 (cited on page 78).

[126] T. Joachims, T. Finley, and C.-N. J. Yu. "Cutting-Plane Training of Structural SVMs". In: *Mach. Learn.* 77.1 (Oct. 2009), pages 27–59 (cited on pages 71, 84, 113).

[127] M. K. Johnson and H. Farid. *Metric measurements on a plane from a single image*. Technical report. Dartmouth College, 2006 (cited on page 87).

[128] M. K. Johnson and H. Farid. "Exposing Digital Forgeries in Complex Lighting Environments". In: *IEEE Trans. Inf. Forensics Security* 2.3 (Sept. 2007), pages 450–461 (cited on page 87).

[129] M. Jung, X. Bresson, T. F. Chan, and L. A. Vese. "Nonlocal Mumford-Shah Regularizers for Color Image Restoration". In: *IEEE Trans. Image Process.* 20.6 (June 2011), pages 1583–1598 (cited on page 53).

[130] P. Kakar and N. Sudha. "Exposing Postprocessed Copy-Paste Forgeries Through Transform-Invariant Features". In: *IEEE Trans. Inf. Forensics Security* 7.3 (June 2012), pages 1018–1028 (cited on page 87).

[131] X. Kang, Y. Li, Z. Qu, and J. Huang. "Enhancing Source Camera Identification Performance With a Camera Reference Phase Sensor Pattern Noise". In: *IEEE Trans. Inf. Forensics Security* 7.2 (2012), pages 393–402 (cited on page 89).

[132] S. Kay. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Prentice Hall, 1993 (cited on page 93).

[133] J. E. Kelley Jr. "The Cutting-Plane Method for Solving Convex Programs". In: *J. Soc. Indust. Appl. Math.* 8.4 (Dec. 1960), pages 703–712 (cited on page 40).

[134] S. Kirkpatrick, C. Gelatt, and M. Vecchi. "Optimization by simulated annealing". In: *Science* 220 (1983), pages 671–680 (cited on page 96).

[135] N. Komodakis and J.-C. Pesquet. "Playing with Duality: An Overview of Recent Primal-Dual Approaches for Solving Large-Scale Optimization Problems". In: *IEEE Signal Process. Mag.* (2015). To appear (cited on pages 4, 6, 8).

[136] A. C. Kot. "Accurate Detection of Demosaicing Regularity for Digital Image Forensics". In: *IEEE Trans. Inf. Forensics Security* 4.4 (Dec. 2009), pages 899–910 (cited on page 88).

[137] B. Krishnapuram, L. Carin, M. A. T. Figueiredo, and A. J. Hartemink. "Sparse Multinomial Logistic Regression: Fast Algorithms and Generalization Bounds". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 27.6 (June 2005) (cited on pages 45, 47, 71, 72, 79).

[138] H. W. Kuhn and A. W. Tucker. "Nonlinear Programming". In: *Proceedings of 2nd Berkeley Symposium*. Berkeley (CA), USA, Oct. 1951, pages 481–492 (cited on page 5).

[139] K. Lang. "Newsweeder: Learning to filter netnews". In: *Proc. of ICML*. Tahoe City, USA, July 1995, pages 331–339 (cited on page 78).

[140] L. Laporte, R. Flamary, S. Canu, S. Déjean, and J. Mothe. "Non-convex Regularizations for Feature Selection in Ranking With Sparse SVM". In: *j-ieee-tnnls* 25.6 (June 2014), pages 1118–1130 (cited on pages 72, 79).

[141] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. "Learning realistic human actions from movies". In: *Proc. of CVPR*. Anchorage, AK, June 2008, pages 1–8 (cited on page 71).

[142] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. "Gradient-Based Learning Applied to Document Recognition". In: *Proc. of IEEE* 86.11 (Nov. 1998), pages 2278–2324 (cited on pages 48, 78).

[143] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. "Gradient-based learning applied to document recognition". In: *Proc. of IEEE* 86.11 (Nov. 1998), pages 2278–2324 (cited on page 78).

[144] Y. LeCun, K. Kavukcuoglu, and C. Farabet. "Convolutional networks and applications in vision". In: *Proc. of ISCAS*. May 2010, pages 253–256 (cited on page 78).

[145] S. Lefkimmiatis, A. Roussos, M. Unser, and P. Maragos. "Convex Generalizations of Total Variation Based on the Structure Tensor with Applications to Inverse Problem". In: *Proc. of SSVM*. Volume 7893. 2013, pages 48–60 (cited on pages 53, 58).

[146] C. t. Li and R. Satta. "An Empirical Investigation into the Correlation between Vignetting Effect and the Quality of Sensor Pattern Noise". In: *IET Comput. Vision* 6.6 (Nov. 2012), pages 560–566 (cited on page 88).

[147] C. T. Li. "Source Camera Identification Using Enhanced Sensor Pattern Noise". In: *IEEE Trans. Inf. Forensics Security* 5.2 (June 2010), pages 280–287 (cited on page 88).

[148] C. T. Li and Y. Li. "Color-Decoupled Photo Response Non-Uniformity for Digital Image Forensics". In: *IEEE Trans. Circuits Syst. Video Technol.* 22.2 (Feb. 2012), pages 260–271 (cited on page 88).

[149] S. Z. Li. *Markov random field modeling in image analysis*. Springer-Verlag, 2001 (cited on page 94).

[150] Y. Li and S. Osher. "A new median formula with applications to PDE based denoising". In: *Commun. Math. Sci.* 7.3 (Sept. 2009), pages 741–753 (cited on page 27).

[151] T. Lin and S. Bourennane. "Survey of hyperspectral image denoising methods based on tensor decompositions". In: *EURASIP J. Adv. Signal Process.* 186 (2013) (cited on page 56).

[152] Z. Lin, J. He, X. Tang, and C. k. Tang. "Fast, automatic and fine-grained tampered JPEG image detection via DCT coefficient analysis". In: *Pattern Recogn.* 42 (2009), pages 2492–2501 (cited on page 87).

[153] B. B. Liu, Y. Hu, and H. K. Lee. "Source camera identification from significant noise residual regions". In: *Proc. of ICIP*. 2010, pages 1749–1752 (cited on page 89).

[154] Q. Liu, X. Cao, C. Deng, and X. Guo. "Identifying image composites through shadow matte consistency". In: *IEEE Trans. Inf. Forensics Security* 6 (2011), pages 1111–1122 (cited on page 87).

[155] Y. Liu, H. Helen Zhang, C. Park, and J. Ahn. "Support Vector Machines with Adaptive Lq Penalty". In: *Comput. Stat. Data Anal.* 51.12 (Aug. 2007), pages 6380–6394 (cited on page 72).

[156] Y. Liu and Y. Wu. "Variable Selection via a Combination of the L0 and L1 Penalties". In: *J. Comp. Graph. Stat.* 14.4 (Dec. 2007), pages 782–798 (cited on page 72).

[157] L. Lorenzi, F. Melgani, and G. Mercier. "Missing-Area Reconstruction in Multispectral Images Under a Compressive Sensing Perspective". In: *IEEE Trans. Geosci. Remote Sens.* 51.7 (July 2013), pages 3998–4008 (cited on page 53).

[158] Y. Lucet. "La transformée de Legendre-Fenchel et la convexifiée d'unefonction: algorithmes rapides de calcul, analyse et régularité du second ordre". PhD thesis. Université Paul Sabatier, Feb. 1997 (cited on page 42).

[159] F. Luisier, T. Blu, and M. Unser. "A New SURE Approach to Image Denoising: Interscale Orthonormal Wavelet Thresholding". In: *IEEE Trans. Image Process.* 16.3 (Mar. 2007), pages 593–606 (cited on page 5).

[160] J. Lukas, J. Fridrich, and M. Goljan. "Detecting digital image forgeries using sensor pattern noise". In: *Proc. of SPIE*. Volume 6072. 2006, pages 362–372 (cited on page 88).

[161] W.-K Ma, J. M. Bioucas-Dias, T.-H. Chan, N. Gillis, P. Gader, et al. "A signal processing perspective on hyperspectral unmixing". In: *IEEE Trans. Signal Process.* 31.1 (Jan. 2014), pages 67–81 (cited on page 54).

[162] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi. "A Nonlocal Transform-Domain Filter for Volumetric Data Denoising and Reconstruction". In: *IEEE Trans. Image Process.* 22.1 (Jan. 2013), pages 119–133 (cited on page 27).

[163] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. "Online Learning for Matrix Factorization and Sparse Coding". In: *J. Mach. Learn. Res.* 11 (Mar. 2010), pages 19–60 (cited on page 27).

[164] S. Mallat. *A wavelet tour of signal processing*. Academic Press, 1997 (cited on page 27).

[165] J. Martin-Herrero. "Anisotropic Diffusion in the Hypercube". In: *IEEE Trans. Geosci. Remote Sens.* 45.5 (May 2007), pages 1386–1398 (cited on page 56).

[166] D. Martín-Iglesias, J. Bernal-Chaves, C. Peláez-Moreno, A. Gallardo-Antolín, and F. Díaz-de María. "A Speech Recognizer Based on Multiclass SVMs with HMM-Guided Segmentation". In: *Nonlinear Analyses and Algorithms for Speech Processing* 3817 (2005), pages 257–266 (cited on page 71).

[167] L. Meier, S. Van De Geer, and P. Bühlmann. "The group Lasso for logistic regression". In: *J. R. Stat. Soc. B* 70.1 (2008), pages 53–71 (cited on page 72).

[168] R. Mendez-Rial, M. Calvino-Cancela, and J. Martin-Herrero. "Accurate implementation of anisotropic diffusion in the hypercube". In: *IEEE Geosci. Remote Sens. Lett.* 7.4 (2010), pages 870–874 (cited on page 56).

[169] K. Miettinen. *Nonlinear Multiobjective Optimization*. Springer, 1998 (cited on pages 4, 5).

[170] M. K. Mihçak, I. Kozintsev, and K. Ramchandran. "Spatially Adaptive Statistical Modeling of Wavelet Image Coefficients and its Application to Denoising". In: *Proc. of ICASSP*. Volume 6. Phoenix, USA, Mar. 1999, pages 3253–3256 (cited on page 90).

[171] T. Miyata and Y. Sakai. "Vectorized total variation defined by weighted L-infinity norm for utilizing inter channel dependency". In: *Proc. of ICIP*. Sept. 2012, pages 3057–3060 (cited on pages 27, 31).

[172] R. Molina, J. Nunez, F. J. Cortijo, and J. Mateos. "Image Restoration in Astronomy: A Bayesian perspective". In: *IEEE Signal Process. Mag.* 18.2 (Mar. 2001), pages 11–29 (cited on page 55).

[173] T. Möllenhoff, E. Strekalovskiy, M. Möller, and D. Cremers. "Low Rank Priors for Color Image Regularization". In: *Proc. of EMMCVPR*. Hong Kong, Jan. 2015 (cited on page 4).

[174] J. J. Moreau. "Proximité et dualité dans un espace hilbertien". In: *Bull. Soc. Math. France* 93 (1965), pages 273–299 (cited on page 5).

[175] G. Obozinski, B. Taskar, and M. I. Jordan. "Joint covariate selection and joint subspace selection for multiple classification problems". In: *Stat. Comput.* 20.2 (2010), pages 231–252 (cited on page 72).

[176] J. Oliveira, M. Figueiredo, and J. M. Bioucas-Dias. "Parametric blur estimation for blind restoration of natural images: linear motion and out-of-focus". In: *IEEE Trans. Image Process.* 23 (2014), pages 466–477 (cited on page 53).

[177] S. Ono and I. Yamada. "Second-order Total Generalized Variation constraint". In: *Proc. of ICASSP*. Florence, Italy, May 2014 (cited on pages 13, 27).

[178] N. Parikh and S. Boyd. "Proximal algorithms". In: *Foundations and Trends in Optimization* 1.3 (2014), pages 123–231 (cited on pages 6, 9, 12, 19, 76).

[179] J.-C. Pesquet, A. Benazza-Benyahia, and C. Chaux. "A SURE approach for digital signal/image deconvolution problems". In: *IEEE Trans. Signal Process.* 57.12 (Dec. 2009), pages 4616–4632 (cited on page 5).

[180] J.-C. Pesquet and N. Pustelnik. "A parallel inertial proximal optimization method". In: *Pac. J. Optim.* 8.2 (Apr. 2012), pages 273–305 (cited on pages 6, 16, 62).

[181] D. W. Peterson. "A Review of Constraint Qualifications in Finite-Dimensional Spaces". In: *SIAM Review* 15.3 (July 1973), pages 639–654 (cited on page 5).

[182] G. Peyré and J. Fadili. "Group Sparsity with Overlapping Partition Functions". In: *Proc. of EUSIPCO*. Barcelona, Spain, Aug. 2011 (cited on page 62).

[183] J. C. Platt. "Fast Training of Support Vector Machines Using Sequential Minimal Optimization". In: *Advances in Kernel Methods*. Edited by B. Schölkopf, C. J. C. Burges, and A. J. Smola. Cambridge, USA: MIT Press, Jan. 1998, pages 185–208 (cited on pages 71, 113).

[184] T. Pock, D. Cremers, H. Bischof, and A. Chambolle. "Global Solutions of Variational Models with Convex Regularization". In: *SIAM J. Imaging Sci.* 3 (2010), pages 1122–1145 (cited on page 4).

[185] A. C. Popescu and H. Farid. "Exposing digital forgeries in color filter array interpolated images". In: *IEEE Trans. Signal Process.* 53.10 (Oct. 2005), pages 3948–3959 (cited on page 88).

[186] N. Pustelnik, C. Chaux, and J.-C. Pesquet. "Parallel ProXimal Algorithm for image restoration using hybrid regularization". In: *IEEE Trans. Image Process.* 20.9 (Sept. 2011), pages 2450–2462 (cited on page 7).

[187] N. Pustelnik, J.-C. Pesquet, and C. Chaux. "Relaxing Tight Frame Condition in Parallel Proximal Methods for Signal Restoration". In: *IEEE Trans. Signal Process.* 60.2 (Feb. 2012), pages 968–973 (cited on page 62).

[188] A. Quattoni, X. Carreras, M. Collins, and T. Darrell. "An Efficient Projection for $\ell_{1,\infty}$ Regularization". In: *Proc. of ICML*. Montreal, Quebec, June 2009 (cited on pages 12, 13, 16, 28, 59, 62).

[189] A. Rakotomamonjy, R. Flamary, G. Gasso, and S. Canu. "$\ell_p - \ell_q$ penalty for sparse linear and sparse multiple kernel multi-task learning". In: *IEEE Trans. Neural Netw.* 22.8 (Aug. 2011), pages 1307–1320 (cited on pages 47, 71).

[190] H. Rauhut, K. Schnass, and P. Vandergheynst. "Compressed Sensing and Redundant Dictionaries". In: *IEEE Trans. Inf. Theory* 54.5 (May 2008), pages 2210–2219 (cited on page 53).

[191] R. Rifkin and A. Klautau. "In defense of one-vs-all classification". In: *Mach. Learn.* 5 (2004), pages 101–141 (cited on page 72).

[192] R. T. Rockafellar and R. J.-B. Wets. *Variational analysis*. Springer-Verlag, 2004 (cited on page 42).

[193] J. Romberg. "Compressive sensing by random convolution". In: *SIAM J. Imaging Sci.* 2.4 (Oct. 2009), pages 1098–1128 (cited on page 53).

[194] L. Rosasco, S. Villa, S. Mosci, M. Santoro, and A. Verri. "Nonparametric Sparsity and Regularization". In: *J. Mach. Learn. Res.* 14 (July 2013), pages 1665–1714 (cited on page 71).

[195] L. Rudin, S. Osher, and E. Fatemi. "Nonlinear Total Variation Based Noise Removal Algorithms". In: *Physica D* 60.1-4 (Nov. 1992), pages 259–268 (cited on pages 4, 15, 26).

[196] G. Sapiro and D. L. Ringach. "Anisotropic Diffusion of Multivalued Images with Applications toColor Filtering". In: *IEEE Trans. Image Process.* 5.11 (Nov. 1996), pages 1582–1586 (cited on page 53).

[197] B. Schölkopf, A. Smola, and K.-R. Müller. "Kernel principal component analysis". In: *Lecture Notes in Computer Science* 1327 (1997), pages 583–588 (cited on pages 48, 78).

[198] S. Setzer, G. Steidl, and T. Teuber. "Deblurring Poissonian images by split Bregman techniques". In: *J. Vis. Commun. Image Represent.* 21.3 (Apr. 2010), pages 193–199 (cited on page 62).

[199] H. Shen, X. Li, L. Zhang, D. Tao, and C. Zeng. "Compressed Sensing-Based Inpainting of Aqua Moderate Resolution Imaging Spectroradiometer Band 6 Using Adaptive Spectrum-Weighted Sparse Bayesian Dictionary Learning". In: *IEEE Trans. Geosci. Remote Sens.* 52.2 (Feb. 2014), pages 894–906 (cited on page 53).

[200] K. Slavakis, I. Yamada, and N. Ogura. "The adaptive projected subgradient method over the fixed point set of strongly attracting nonexpansive mappings". In: *Numer. Func. Anal. Opt.* 27.7-8 (Nov. 2006), pages 905–930 (cited on page 11).

[201] N. Sochen, R. Kimmel, and R. Malladi. "A General Framework for Low Level Vision". In: *IEEE Trans. Image Process.* 7.3 (Mar. 1998), pages 310–318 (cited on page 53).

[202] G. W. Stewart. "On the Early History of the Singular Value Decomposition". In: *SIAM Review* 35.4 (1993), pages 551–566 (cited on pages 57, 60).

[203] E. Strekalovskiy and D. Cremers. "Real-Time Minimization of the Piecewise Smooth Mumford-Shah Functional". In: *Proc. of ECCV*. Zürich, Switzerland, Sept. 2014, pages 127–141 (cited on page 4).

[204] C. Studer, T. Goldstein, W. Yin, and R. G. Baraniuk. "Democratic Representations". In: *preprint* (Jan. 2014). http://arxiv.org/abs/1401.3420 (cited on page 16).

[205] C. Sutour, C.-A. Deledalle, and J.-F. Aujol. "Adaptive regularization of the NL-means: Application to image and video denoising". In: *IEEE Trans. Image Process.* 23.8 (Aug. 2014), pages 3506–3521 (cited on page 27).

[206] M. Tan, L. Wang, and I. W. Tsang. "Learning sparse SVM for feature selection on very high dimensional datasets". In: *Proc. of ICML*. Haifa, Israel, June 2010, pages 1047–1054 (cited on page 72).

[207] T. Teuber, G. Steidl, and R. H. Chan. "Minimization and Parameter Estimation for Seminorm RegularizationModels with I-Divergence Constraints". In: *Inverse Probl.* 29.3 (2013), page 035007 (cited on pages xxviii, 4, 5, 12, 13, 45, 46).

[208] R. Tibshirani. "Regression Shrinkage and Selection via the Lasso". In: *J. R. Statist. Soc. B* 58.1 (1996), pages 267–288 (cited on page 4).

[209] A. Tikhonov. "Tikhonov regularization of incorrectly posed problems". In: *Soviet Math. Dokl.* 4 (1963), pages 1624–1627 (cited on pages 3, 5, 15, 26).

[210] D. Tschumperlé and R. Deriche. "Constrained and Unconstrained PDE's for Vector Image Restoration". In: *Proc. SCIA*. Bergen, Norway, June 2001 (cited on page 53).

[211] D. Tschumperlé and R. Deriche. "Vector-Valued Image Regularization with PDEs: A Common Framework for Different Applications". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 27.4 (Apr. 2005), pages 506–517 (cited on page 56).

[212] P. Tseng. "Convergence of a block coordinate descent method for nondifferentiable minimization". In: *J. Optim. Theory Appl.* 109.3 (June 2001), pages 475–494 (cited on page 5).

[213] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. "Large Margin methods for structured and interdependent output variables". In: *J. Mach. Learn. Res.* 6 (2005), pages 1453–1484 (cited on pages 71, 113).

[214] D. Tuia, M. Volpi, M. Dalla Mura, A. Rakotomamonjy, and R. Flamary. "Automatic Feature Learning for Spatio-Spectral Image Classification With Sparse SVM". In: *IEEE Trans. Geosci. Remote Sens.* 52.10 (Oct. 2014), pages 6062–6074 (cited on pages 47, 71).

[215] B. C. Vũ. "A splitting algorithm for dual monotone inclusions involving cocoercive operators". English. In: *Adv. Comput. Math.* 38.3 (Apr. 2013), pages 667–681 (cited on pages 6, 8, 45, 49).

[216] P. Valero S. Salembier and J. Chanussot. "Hyperspectral Image Representation and Processing With Binary Partition Trees". In: *IEEE Trans. Image Process.* 22.4 (Apr. 2013), pages 1430–1443 (cited on page 55).

[217] E. Van Den Berg and M. P. Friedlander. "Probing the Pareto frontier for basis pursuit solutions". In: *SIAM J. Sci. Comput.* 31.2 (Nov. 2008), pages 890–912 (cited on pages 4, 5, 10, 12, 13, 28, 59, 62, 67, 97).

[218] S. Villa, L. Rosasco, S. Mosci, and A. Verri. "Proximal methods for the latent group lasso penalty". In: *Comput. Optim. Appl.* 58.2 (Dec. 2014), pages 381–407 (cited on pages 47, 71).

[219] C. Vonesch, F. Aguet, J.-L. Vonesch, and M. Unser. "The Colored Revolution of Bioimaging". In: *IEEE Signal Process. Mag.* 23.3 (May 2006), pages 20–31 (cited on page 55).

[220] L. Wang and X. Shen. "On l_1-norm multi-class support vector machines: methodology and theory". In: *J. Am. Statist. Assoc.* 102 (2007), pages 583–594 (cited on pages 72, 82).

[221] L. Wang, J. Zhu, and H. Zou. "The doubly regularized support vector machine". In: *Stat. Sinica* 16 (2006), pages 589–616 (cited on page 72).

[222] J. Weickert. "Coherence-enhancing diffusion of colour images". In: *Image Vision Comput.* 17.3-4 (Mar. 1999), pages 201–212 (cited on page 53).

[223] P. Weiss, L. Blanc-Féraud, and G. Aubert. "Efficient Schemes for Total Variation Minimization Under Constraints in Image Processing". In: *SIAM J. Sci. Comput.* 31.3 (Apr. 2009), pages 2047–2080 (cited on pages 59, 62).

[224] J. Weston, A. Elisseeff, B. Schölkopf, and M. Tipping. "Use of the Zero-Norm With Linear Models and Kernel Methods". In: *Mach. Learn.* 3 (2002), pages 1439–1461 (cited on page 72).

[225] S. J. Wright. *Primal-dual interior-point methods*. SIAM, 1997 (cited on page 5).

[226] J. Wu, F. Liu, L. C. Jiao, X. Wang, and B. Hou. "Multivariate Compressive Sensing for Image Reconstruction in the Wavelet Domain". In: *IEEE Trans. Image Process.* 20.12 (Dec. 2011), pages 3483–3494 (cited on page 15).

[227] L. Xu and J. Jia. "Two-Phase Kernel Estimation for Robust Motion Deblurring". In: *Proc. of ECCV*. Crete, Greece, Sept. 2010 (cited on page 27).

[228] L. Xu, C. Lu, Y. Xu, and J. Jia. "Image Smoothing via L0 Gradient Minimization". In: *ACM T. Graphics* 30.6 (2011) (cited on page 27).

[229] Z. Yang and M. Jacob. "Nonlocal regularization of inverse problems: a unified variational framework". In: *IEEE Trans. Image Process.* 22.8 (Aug. 2013), pages 3192–3203 (cited on page 27).

[230] J. S. Yedidia, W. T. Freeman, and Y. Weiss. "Generalized belief propagation". In: *Proc. of NIPS* (2000), pages 689–695 (cited on page 96).

[231] I. Yerushalmy and H. Hel-Or. "Digital Image Forgery Detection Based on Lens and Sensor Aberration". In: *Int. J. Comput. Vision* 92.1 (Nov. 2011), pages 71–91 (cited on page 88).

[232] D. C. Youla and H. Webb. "Image restoration by the method of convex projections. Part I - Theory". In: *IEEE Trans. Med. Imag.* 1.2 (Oct. 1982), pages 81–94 (cited on page 5).

[233] G.-X. Yuan, K.-W. Chang, C.-J. Hsieh, and C.-J. Lin. "A Comparison of Optimization Methods and Software for Large-scale L1-regularized Linear Classification". In: *Mach. Learn.* 11 (Dec. 2010), pages 3183–3234 (cited on page 71).

[234] M. Yuan and Y. Lin. "Model selection and estimation in regression with grouped variables". In: *J. R. Stat. Soc. B* 68 (2006), pages 49–67 (cited on page 72).

[235] Q. Yuan, L. Zhang, and H. Shen. "Hyperspectral Image Denoising Employing a Spectral-Spatial Adaptive Total Variation Model". In: *IEEE Trans. Geosci. Remote Sens.* 50.10 (Oct. 2012), pages 3660–3677 (cited on pages xxxiii, 56, 58, 63, 65, 66).

[236] C. Zach, T. Pock, and H. Bischof. "A duality based approach for real-time TV-L1 optical flow". In: *DAGM Symposium*. Heidelberg, Germany, Sept. 2007, pages 214–223 (cited on page 53).

[237] F. Zach, C. Riess, and E. Angelopoulou. "Automated Image Forgery Detection through Classification of JPEG Ghosts". In: *Pattern Recogn.* 7476 (2012), pages 185–194 (cited on page 87).

[238] R. Zanella, P. Boccacci, L. Zanni, and M. Bertero. "Efficient gradient projection methods for edge-preserving removal of Poisson noise". In: *Inverse Probl.* 25.4 (2009) (cited on page 45).

[239] C. Zhang and H. Zhang. "Exposing Digital Image Forgeries by Using Canonical Correlation Analysis". In: *Proc. of ICPR*. 2010, pages 838–841 (cited on page 89).

[240] H. Zhang, L. Zhang, and H. Shen. "A super-resolution reconstruction algorithm for hyperspectral images". In: *Signal Process.* 92.9 (Sept. 2012), pages 2082–2096 (cited on pages 53, 54, 56, 58, 112).

[241] H. Zhang, Y. Liu, Y. Wu, and J. Zhu. "Variable selection for multicategory SVM via sup-norm regularization". In: *Electron. J. Stat.* 2 (2008), pages 149–167 (cited on pages 72, 82).

[242] T. Zhang. "Solving Large Scale Linear Prediction Problems Using Stochastic Gradient Descent Algorithms". In: *Proc. of ICML*. Banff, Canada, July 2004, pages 919–926 (cited on pages 75, 113).

[243] Y. Zhao, S. Wang, X. Zhang, and H. Yao. "Robust Hashing for Image Authentication Using Zernike Moments and Local Features". In: *IEEE Trans. Inf. Forensics Security* 8.1 (Jan. 2013), pages 55–63 (cited on page 87).

[244] G. Zhou and D. Lv. "An Overview of Digital Watermarking in Image Forensics". In: *Int. Conf. Comput. Sciences Optim.* Apr. 2011, pages 332–335 (cited on page 87).

[245] H. Zou and M. Yuan. "The f_∞-norm support vector machine". In: *Stat. Sinica* 18 (2008), pages 379–398 (cited on page 72).

Éclatement épigraphique de contraintes convexes. Application à la restauration d'images, la classification supervisée, et la détection d'images falsifiées.

Giovanni CHIERCHIA

RESUME : Dans cette thèse, nous proposons une approche d'optimisation convexe pour aborder des problèmes en restauration d'images multi-composantes, en apprentissage supervisé et en détection d'images falsifiées. Le fil conducteur de ces problèmes est la présence de contraintes convexes non linéaires qui sont difficiles à gérer avec les méthodes de l'état-de-l'art. Par conséquent, nous avons élaboré une technique d'éclatement épigraphique pour en simplifier la gestion. En s'appuyant sur cette approche, nous avons également proposé des contributions spécifiques pour les applications susmentionnées.

MOTS-CLEFS : Optimisation convexe, méthodes proximales, contraintes non linéaires, restauration d'images multi-composantes, apprentissage supervisé, détection d'images falsifiées.

ABSTRACT : In this thesis, we present a convex optimization approach to address three problems arising in multicomponent image recovery, supervised classification, and image forgery detection. The common thread among these problems is the presence of nonlinear convex constraints difficult to handle with state-of-the-art methods. Therefore, we present a splitting technique to simplify the management of such constraints. Relying on this approach, we also propose some contributions to the aforementioned applications.

KEY-WORDS : Convex optimization, proximal methods, nonlinear constraints, multicomponent image restoration, supervised classification, image forgery detection.

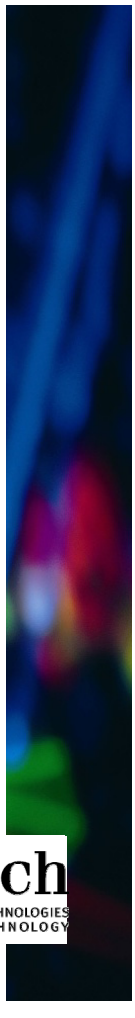

TELECOM ParisTech

EDiTE

ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY