



HAL
open science

Submodular Functions: from Discrete to Continuous Domains

Francis Bach

► **To cite this version:**

| Francis Bach. Submodular Functions: from Discrete to Continuous Domains. 2015. hal-01222319v1

HAL Id: hal-01222319

<https://hal.science/hal-01222319v1>

Preprint submitted on 29 Oct 2015 (v1), last revised 23 Feb 2016 (v2)

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

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

Submodular Functions: from Discrete to Continuous Domains

Francis Bach
INRIA - Sierra project-team
Département d'Informatique de l'Ecole Normale Supérieure
Paris, France
`francis.bach@ens.fr`

October 29, 2015

Abstract

Submodular set-functions have many applications in combinatorial optimization, as they can be minimized and approximately maximized in polynomial time. A key element in many of the algorithms and analyses is the possibility of extending the submodular set-function to a convex function, which opens up tools from convex optimization. Submodularity goes beyond set-functions and has naturally been considered for problems with multiple labels or for functions defined on continuous domains, where it corresponds essentially to cross second-derivatives being nonpositive. In this paper, we show that most results relating submodularity and convexity for set-functions can be extended to all submodular functions. In particular, (a) we naturally define a continuous extension in a set of measures, (b) show that the extension is convex if and only if the original function is submodular, (c) prove that the problem of minimizing a submodular function is equivalent to a typically non-smooth convex optimization problem, and (d) propose another convex optimization problem with better computational properties (e.g., a smooth dual problem). Most of these extensions from the set-function situation are obtained by drawing links with the theory of multi-marginal optimal transport. We then provide practical algorithms to minimize generic submodular functions on discrete domains, with associated convergence rates.

1 Introduction

Submodularity has emerged as an important concept in combinatorial optimization, akin to convexity in continuous optimization, with many applications in machine learning, computer vision or signal processing [1, 2, 3, 4]. Most of the literature on submodular functions focuses on *set-functions*, i.e., functions defined on the set of subsets of a given base set. Such functions are classically equivalently obtained as functions defined on the vertices of the hypercube $\{0, 1\}^n$, if n is the cardinality of the base set. Throughout the paper, we will make this identification and refer to *set-functions* as functions defined on $\{0, 1\}^n$.

Like convex functions, submodular set-functions can be minimized exactly in polynomial time, by drawing links with convexity, leading to either combinatorial algorithms akin to max-flow algorithms [5, 6, 7], or algorithms based on a convex optimization techniques [4, Section 10]. Unlike convex functions, submodular set-functions can also be *maximized* approximately in polynomial time with simple greedy algorithms that come with approximation guarantees [8, 9].

In this paper, we focus primarily on submodular function minimization and links with convexity. In the set-function situation, it is known that submodular function minimization is equivalent to

a convex optimization problem, which is obtained by considering a *continuous extension* of the submodular function, from vertices of the hypercube $\{0, 1\}^n$ to the full hypercube $[0, 1]^n$. This extension, usually referred to as the Choquet integral [10] or the Lovász extension [11], is convex if and only if the original set-function is submodular. Moreover, when the set-function is submodular, minimizing the original set-function or the convex extension is equivalent. Finally, simple and efficient algorithms based on generic convex optimization algorithms may be used for minimization.

The main goal of this paper is to show that all of these results naturally extend to all submodular functions defined more generally on subsets of \mathbb{R}^n . In this paper, we focus on functions defined on subsets of \mathbb{R}^n of the form $\mathcal{X} = \prod_{i=1}^n \mathcal{X}_i$, where each $\mathcal{X}_i \subset \mathbb{R}$. A function $H : \mathcal{X} \rightarrow \mathbb{R}$, is then submodular if and only if for all $(x, y) \in \mathcal{X} \times \mathcal{X}$,

$$H(x) + H(y) \geq H(\max\{x, y\}) + H(\min\{x, y\}),$$

where the min and max operations are applied component-wise. This extended notion of submodularity has been thoroughly studied [12, 13], and some of the results on submodular set-functions (which correspond to $\mathcal{X}_i = \{0, 1\}$ for all $i \in \{1, \dots, n\}$) have already been extended, such as the possibility of minimizing discrete functions (i.e., when all \mathcal{X}_i are finite) in polynomial time. This is done usually by a reduction to the problem of minimizing a submodular function on a ring family [5]. Moreover, particular examples, such as certain cuts with ordered labels [14, 15, 16] leads to min-cut/max-flow reformulations. Finally, for functions defined as sums of local functions, it is known that the usual linear programming relaxations are tight for these submodular functions (see [17] and references therein).

While finite sets already lead to interesting applications in computer vision [14, 15, 16], functions defined on products of sub-intervals of \mathbb{R} are particularly interesting. Indeed, when twice-differentiable, a function is submodular if and only if all cross-second-order derivatives are non-positive, i.e., for all $x \in \mathcal{X}$:

$$\frac{\partial^2 H}{\partial x_i \partial x_j}(x) \leq 0.$$

In this paper, we provide simple algorithms based solely on function evaluations to minimize all of these functions. This thus opens up a new set of “simple” functions that can be efficiently minimized, which neither is included nor includes convex functions, with potentially many interesting theoretical or algorithmic developments.

In this paper, we make the following contributions:

- We propose in Section 2.2 a continuous extension in a set of measures and show in Section 2.3 that it is convex if and only if the function is submodular and that for submodular functions, minimizing the extension, which is a convex optimization problem, is equivalent to minimizing the original function. This is made by drawing links with the theory of optimal transport, which provides simple intuitive proofs (even for submodular set-functions).
- We show in Section 2.5 that minimizing the extension plus a well-defined separable convex function is equivalent to minimizing a series of submodular functions.
- For finite sets, we show in Section 3 a direct link with existing notions for submodular set-functions, such as the base polytope. In the general situation, two polyhedra naturally emerge (instead of one).
- For finite sets, we provide two sets of algorithms for minimizing submodular functions, one based on a non-smooth optimization problem on a set of measures (projected subgradient descent), and one based on smooth functions and “Frank-Wolfe” methods. They can be readily applied to all situations by discretizing the sets if continuous. We provide in Section 5 a simple experiment on a one-dimensional signal denoising problem.

In this paper, we assume basic knowledge of convex analysis (see, e.g., [18, 19]) and submodular analysis (see, e.g., [20, 4]).

2 Submodular functions

We consider a measurable function $H : \mathcal{X} = \prod_{i=1}^n \mathcal{X}_i \rightarrow \mathbb{R}$, defined on the product of n sets \mathcal{X}_i , which are included in \mathbb{R} , and thus equipped with a total order. Typically, \mathcal{X}_i will be a finite set such as $\{0, \dots, k_i - 1\}$ or an interval (which we refer to as a continuous domain).

The function H is said to be *submodular* if and only if [12, 13]:

$$\forall (x, y) \in \mathcal{X} \times \mathcal{X}, \quad H(x) + H(y) \geq H(\max\{x, y\}) + H(\min\{x, y\}), \quad (1)$$

where the min and max operations are applied component-wise. An equivalent definition is that for any $x \in \mathcal{X}$ and (different) basis vectors e_i, e_j and $a_i, a_j \in \mathbb{R}_+$ such that $x_i + a_i \in \mathcal{X}_i$ and $x_j + a_j \in \mathcal{X}_j$, then

$$H(x + a_i e_i) + H(x + a_j e_j) \geq H(x) + H(x + a_i e_i + a_j e_j). \quad (2)$$

Moreover, we only need the statement above in the limit of a_i and a_j tending to zero (but different from zero and such that $x_i + a_i \in \mathcal{X}_i$ and $x_j + a_j \in \mathcal{X}_j$).

Submodularity-preserving operations. Like for set-functions, the set of submodular functions is a cone. Unlike set-functions, the notion of restricting the set \mathcal{X}_i is non-trivial: any function defined by restriction on a product of subsets of \mathcal{X}_i is submodular. This will be useful when discretizing a continuous domain in Section 4.

Moreover, submodularity is invariant by separable increasing reparameterizations, that is, if for all $i \in \{1, \dots, n\}$, $\varphi_i : \mathcal{X}_i \rightarrow \mathcal{X}_i$ is an increasing bijection, H is submodular, if and only if, $x \mapsto H[\varphi_1(x_1), \dots, \varphi_n(x_n)]$ is submodular. Note the difference with convex functions which are invariant by affine reparameterizations.

Strict submodularity. We define the notion of strict submodularity through a strict inequality in Eq. (1) for any two x and y which are not comparable [13].

Modular functions. We define modular functions as functions H such that both H and $-H$ are submodular. This happens to be equivalent to H being a *separable* function, that is a function that is a sum of n functions that depend arbitrarily on single variables [13, Theorem 3.3]. These will play the role that linear functions play for convex functions. Note that when each \mathcal{X}_i is a sub-interval of \mathbb{R} , this set of functions is much larger than the set of linear functions.

2.1 Examples

In this section, we provide simple examples of submodular functions. We will use as running examples throughout the paper submodular set-functions (to show that our new results directly extend the ones for set-functions), modular functions (because they provide a very simple example of the concepts we introduce), and functions that are sums of terms $\varphi_{ij}(x_i - x_j)$ where φ_{ij} is convex (for the link with Wasserstein distances between probability measures).

- **Set-functions:** When each \mathcal{X}_i has exactly two elements, e.g., $\mathcal{X}_i = \{0, 1\}$, we recover exactly submodular set-functions defined on $\{1, \dots, n\}$, with the usual identification of $\{0, 1\}^n$ with the set of subsets of $\{1, \dots, n\}$.
- **Functions on intervals:** When each \mathcal{X}_i is an interval of \mathbb{R} and H is twice differentiable on \mathcal{X} , then H is submodular if and only if all cross-second-derivatives are non-negative, i.e.,

$$\forall i \neq j, \forall x \in \mathcal{X}, \frac{\partial^2 H}{\partial x_i \partial x_j}(x) \leq 0.$$

This includes functions of the form $\varphi_{ij}(x_i - x_j)$ for $\varphi_{ij} : \mathbb{R} \rightarrow \mathbb{R}$ convex, and $x \mapsto g(\sum_{i=1}^n \lambda_i x_i)$ for g concave and (λ_i) non-negative weights; this gives examples of functions which are submodular, but convex or concave.

Other examples are $\varphi(\sum_{i=1}^n \lambda_i x_i) - \sum_{i=1}^n \lambda_i \varphi(x_i)$ for φ strictly concave and $\lambda \in \mathbb{R}^n$ in the interior of the simplex, which is non-negative and zero if and only if all x_i are equal.

Moreover, functions of the form $x \mapsto \log \det(\sum_{i=1}^n x_i A_i)$, where A_i are positive definite matrices and $x \geq 0$, are submodular—this extends to other spectral functions [21]. Moreover, if g is the Lovász extension of a submodular set-function, then it is submodular (as a function defined on \mathbb{R}^n)—see proof in Appendix A.1. These give examples of functions which are both convex and submodular. Similarly, the multi-linear extension of a submodular set-function [22], defined on $[0, 1]^n$, is submodular as soon as the original set-function is submodular (see Appendix A.2).

For algorithms, these functions will be approximated on a discrete grid (one separate grid per variable, with a total complexity which is linear in the dimension n), but most of our formulations and convex analysis results extend in the continuous setting with appropriate regularity assumptions. In order to make the exposition as simple as possible, we do not provide details in this paper.

- **Discrete labels:** in this paper, we will often consider the case where the sets \mathcal{X}_i are all finite. They will serve as approximations of functions defined on intervals. We will still use a functional notation to make the extension to continuous settings explicit.

Submodular minimization problems. In this paper, we focus on simple and efficient methods to minimize general submodular functions, based only on function evaluations. Many examples come from signal and image processing, with functions to minimize of the form

$$H(x) = \sum_{i=1}^n f_i(x_i) + \sum_{C \in \mathcal{C}} f_C(x_C),$$

where \mathcal{C} is a set of small subsets (often a set of edges) and each f_C is submodular (while each f_i may be arbitrary) [14, 15, 16]. We consider a simple one-dimensional example in Section 5 as an illustration.

Another motivating example is a probabilistic modelling problem where submodularity on continuous domains appears naturally, namely probabilistic models on $\{0, 1\}^n$ with log-densities which are negatives of submodular functions [23, 24], that is $p(x) = \frac{1}{Z} \exp(-F(x))$, with F submodular and Z the normalizing constant equal to $Z = \sum_{x \in \{0, 1\}^n} \exp(-F(x))$, which is typically hard to compute. In this context, *mean field* inference aims at approximating p by a product of independent distributions $\mu(x) = \prod_{i=1}^n \mu_i(x_i)$, by minimizing the Kullback-Leibler divergence between μ and p , that is, by minimizing

$$\sum_{x \in \mathcal{X}} \mu(x) \log \frac{\mu(x)}{p(x)} = \sum_{i=1}^n \{ \mu_i(1) \log \mu_i(1) + \mu_i(0) \log \mu_i(0) \} + \sum_{x \in \mathcal{X}} \mu(x) F(x) + Z.$$

The first term in the right-hand side is separable, hence submodular, while the second term is exactly the multi-linear extension of the submodular set-function F , which is itself a submodular function (see [22] and Appendix A.2). This implies that in this context, mean field inference may be done globally with arbitrary precision in polynomial time. Note that in this context, previous work [23] has considered replacing the multi-linear extension by the Lovász extension, which is also submodular but also convex (it then turns out to correspond to a different divergence between μ and q).

2.2 Extensions

In this section, in order to avoid fine topological issues, we assume that the sets \mathcal{X}_i are all finite, with k_i elements. More precisely, where we need to enumerate the elements of \mathcal{X}_i , we will consider it to be equal to $\{0, \dots, k_i - 1\}$. All our results extend with appropriate measurability and integrability assumptions (in fact, we will try to use expressions that do not involve k_i most of the time). Note that since H is a function on a discrete set, it is bounded and all integrals are thus well-defined.

Our goal is to minimize the function H through a convex relaxation. Since all our sets \mathcal{X}_i are subsets of \mathbb{R} , we could look for extensions to \mathbb{R}^n directly; this in fact exactly the approach for functions defined on $\{0, 1\}^n$. However, when the sets \mathcal{X}_i have more than two elements, we are going to consider the convex set $\mathcal{P}(\mathcal{X}_i)$ of *probability measures* μ_i on \mathcal{X}_i , which is essentially the convex hull of all Dirac measures; for $\mathcal{X}_i = \{0, \dots, k_i - 1\}$, this is essentially a simplex in dimension k_i . In order to get an *extension*, we look for a function defined on $\mu \in \mathcal{P}^\otimes(\mathcal{X}) = \prod_{i=1}^n \mathcal{P}(\mathcal{X}_i)$, such that if all μ_i , $i = 1, \dots, n$, are Dirac measures at points $x_i \in \mathcal{X}_i$, then we have a function value equal to $H(x_1, \dots, x_n)$.

For a distribution $\mu_i \in \mathcal{P}(\mathcal{X}_i)$ defined on a totally ordered set \mathcal{X}_i , we can define the (reversed) cumulative distribution function $F_{\mu_i} : \mathcal{X}_i \rightarrow [0, 1]$ as $F_{\mu_i}(x_i) = \mu_i(\{y_i \in \mathcal{X}_i, y_i \geq x_i\})$. This is a non-increasing left-continuous function from \mathcal{X}_i to $[0, 1]$. See an illustration in Figure 1 (left). When \mathcal{X}_i is assumed discrete with k_i elements, it may be exactly represented as a vector in \mathbb{R}^{k_i-1} elements with non-decreasing components, that is, given μ_i , we define $\rho_i(x_i) = \mu_i(x_i) + \dots + \mu_i(k_i - 1) = F_{\mu_i}(x_i)$, for $x_i \in \{1, \dots, k_i - 1\}$. Because the measure μ_i has unit total mass, $\rho_i(0)$ is always equal to 1 and can thus be omitted. For example, for $k_i = 2$ (and $\mathcal{X}_i = \{0, 1\}$), then we simply have $\rho_i \in [0, 1]$ which represents the probability that the associated random variable is equal to 1.

Note that in order to preserve the parallel with submodular set-functions, we choose to deviate from the original definition of the cumulative function by considering the mass of the set $\{y_i \in \mathcal{X}_i, y_i \geq x_i\}$ (and not the other direction).

We can define the “inverse” cumulative function from $[0, 1]$ to \mathcal{X}_i as

$$F_{\mu_i}^{-1}(t) = \sup\{x_i \in \mathcal{X}_i, F_{\mu_i}(x_i) \leq t\}.$$

The function $F_{\mu_i}^{-1}$ is non-increasing and right-continuous; we have $F_{\mu_i}(x_i) \leq t_i \Leftrightarrow F_{\mu_i}^{-1}(t_i) \leq x_i$. See an illustration in Figure 1 (right). When \mathcal{X}_i is assumed discrete, $F_{\mu_i}^{-1}$ is piecewise constant with steps at every t equal to $F_{\mu_i}(x_i)$ for a certain x_i . For $k_i = 2$, we get $F_{\mu_i}^{-1}(t) = 1$ if $t < \mu_i(1)$ and 0 if $t > \mu_i(1)$. What happens at $t = \mu_i(1)$ does not matter because this corresponds to a set of zero Lebesgue measure.

We define the extension of H to $\mathcal{P}^\otimes(\mathcal{X}) = \prod_{i=1}^n \mathcal{P}(\mathcal{X}_i)$ as follows:

$$\forall \mu \in \prod_{i=1}^n \mathcal{P}(\mathcal{X}_i), \quad h(\mu_1, \dots, \mu_n) = \int_0^1 H[F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)] dt. \quad (3)$$




Figure 1: Left: cumulative function for a distribution on the real line supported in the set $\{0, \dots, k_i - 1\}$ (note that we could also restrict the domain to be exactly $\{0, \dots, k_i - 1\}$, and the cumulative function would only be composed of the set of plain dots). Right: inverse cumulative function (which would be the same for the distribution with discrete domain).

If all μ_i , $i = 1, \dots, n$ are Diracs at $x_i \in \mathcal{X}_i$, then for all $t \in (0, 1)$, $F_{\mu_i}^{-1}(t) = x_i$ and we indeed have the extension property (again, what happens for $t = 0$ or $t = 1$ is irrelevant because this is on a set of zero Lebesgue measure). For $\mathcal{X}_i = \{0, 1\}$ for all i , then the extension is defined on $[0, 1]^n$ and is equal to $h(\mu) = \int_0^1 H_{\text{set}}(1_{\{\mu(1) \geq t\}}) dt$ and we exactly recover the Choquet integral (i.e., the Lovász extension) for set-functions.

Note that the extension is defined on all tuples of measures (μ_i) but it can equivalently be applied to non-increasing functions from \mathcal{X}_i to $[0, 1]$. As we will see for discrete domains in Section 3, it may also be defined for all non-increasing functions with no constraints to be in $[0, 1]$.

Examples. For our three running examples, we may look at the extension. For set-functions, we recover the usual Choquet integral; for modular functions $H(x) = \sum_{i=1}^n H_i(x_i)$, then we have $h(\mu) = \sum_{i=1}^n \int_{\mathcal{X}_i} H_i(x_i) d\mu_i(x_i)$ which is the expectation of $H(x)$ under the product measure defined by μ . Finally, for the function $\varphi_{ij}(x_i - x_j)$, we obtain a Wasserstein distance between the measures μ_i and μ_j (which is a distance between their cumulative functions) [25].

2.3 Relationship with convexity

Assume each \mathcal{X}_i is finite, $i = 1, \dots, n$. We now prove our first result relating convexity and submodularity, that extends the similar result of [11] from set-functions to all functions. The proof is based on drawing explicit links with the theory of optimal transport; while an implicit connection had been made for $n = 2$ through monotonicity properties of optimal assignment problems [26], the link we propose here is novel.

Theorem 1 (Convexity and submodularity) *Assume each \mathcal{X}_i is finite, $i = 1, \dots, n$. The extension h defined in Eq. (3) is convex if and only if H is submodular.*

Proof We follow a different more direct proof than for submodular set-functions. Given μ , we define an optimal transport problem:

$$\tilde{h}(\mu_1, \dots, \mu_n) = \inf_{\gamma \in \Pi(\mu)} \int_{\mathcal{X}} H(x) d\gamma(x), \quad (4)$$

over all probability measures γ on \mathcal{X} such that the i -th marginal is equal to μ_i . This is the Monge-Kantorovich formulation of the *multi-marginal optimal transport problem* [27].

We first assume that H is submodular. Optimal transport problems on subsets of real numbers with submodular costs are known to have closed-form solutions [28], among which the joint distribution γ such that

$$\gamma\left(\prod_{i=1}^n \{x_i \in \mathcal{X}_i, x_i \geq t_i\}\right) = \max_{i \in \{1, \dots, n\}} F_{\mu_i}(t_i),$$

leading to a closed-form expression for the value of the optimal transport problem:

$$\tilde{h}(\mu) = \int_0^1 H[F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)] dt,$$

which is exactly equal to $h(\mu)$ for all $\mu \in \mathcal{P}^{\otimes}(\mathcal{X})$. Since the optimal transport problem is jointly convex in μ and γ , the resulting minimum $h(\mu)$ is convex in μ .

We now assume that the function h is convex. Following the argument of [12] for the related problem of rearrangement inequalities and [11] for submodular set-functions, we consider two arbitrary elements a and b in \mathcal{X} , and the Dirac measures δ_{a_i} and δ_{b_i} . We have, by convexity:

$$h\left(\frac{1}{2}\delta_{a_1} + \frac{1}{2}\delta_{b_1}, \dots, \frac{1}{2}\delta_{a_n} + \frac{1}{2}\delta_{b_n}\right) \leq \frac{1}{2}h(\delta_{a_1}, \dots, \delta_{a_n}) + \frac{1}{2}h(\delta_{b_1}, \dots, \delta_{b_n}).$$

The right-hand side is equal to $\frac{1}{2}H(a) + \frac{1}{2}H(b)$, while we can compute the left-hand side by computing $F_{\frac{1}{2}\delta_{a_i} + \frac{1}{2}\delta_{b_i}}(x_i)$, which is equal to 0 if $x_i > \max\{a_i, b_i\}$, 1 if $x_i < \min\{a_i, b_i\}$ and $\frac{1}{2}$ if $x_i \in (\min\{a_i, b_i\}, \max\{a_i, b_i\})$. This implies that $F_{\frac{1}{2}\delta_{a_i} + \frac{1}{2}\delta_{b_i}}^{-1}(t)$ is equal to $\min\{a_i, b_i\}$ if $t > \frac{1}{2}$, and to $\max\{a_i, b_i\}$ if $t < \frac{1}{2}$. Thus the left-hand side of the inequality above is $\frac{1}{2}H(\min\{a, b\}) + \frac{1}{2}H(\max\{a, b\})$. Hence, the submodularity. \blacksquare

Like for submodular set-functions [4, Section 5.1], the formulation in Eq. (4) may be seen as a *convex closure* operation, and \tilde{h} should then be an extension of H , and minimizing \tilde{h} and H should be equivalent (as shown below). The key property that submodularity brings is that \tilde{h} may be computed in “closed form”, i.e., without solving an optimization problem.

We now show that minimizing the extension is equivalent to minimizing the original function, implying that we may minimize any submodular function as a convex optimization problem:

Theorem 2 (Equivalent minimization problems) *Assume each \mathcal{X}_i is finite, $i = 1, \dots, n$. If H is submodular, then*

$$\inf_{x \in \mathcal{X}} H(x) = \inf_{\mu \in \mathcal{P}^{\otimes}(\mathcal{X})} h(\mu). \quad (5)$$

Proof Since h is the convex closure of H , the two infima have to be equal. Indeed, we have

$$\inf_{\mu \in \mathcal{P}^{\otimes}(\mathcal{X})} h(\mu) = \inf_{\mu \in \mathcal{P}^{\otimes}(\mathcal{X})} \inf_{\gamma \in \Pi(\mu)} \int_{\mathcal{X}} H(x) d\gamma(x),$$

which is the infimum over all probability measures on \mathcal{X} (without any marginal constraints). It is thus achieved at a Dirac measure at any minimizer $x \in \mathcal{X}$ of $H(x)$.

Moreover, given a minimizer μ for the convex problem,

$$\inf_{x \in \mathcal{X}} H(x) = h(\mu) = \int_0^1 H[F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)] dt \geq \int_0^1 \inf_{x \in \mathcal{X}} H(x) dt = \inf_{x \in \mathcal{X}} H(x).$$

Thus, for almost all $t \in [0, 1]$, $(F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)) \in \mathcal{X}$ is a minimizer of H . \blacksquare

From the proof, we see that a minimizer of H may be obtained from a minimizer of h . Many properties are known regarding minimizers of submodular functions [13], i.e., if x and y are minimizers of H , so are $\min\{x, y\}$ and $\max\{x, y\}$. Moreover, when H is strictly submodular, then the set of minimizers forms a chain.

In practice, given μ and for discrete domains, we can look at the minimal value of $H[F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)]$ over all $t \in [0, 1]$ by enumerating and sorting all possible values of $F_{\mu_i}(x_i)$ (see Section 3.1).

2.4 General dual representations

Given the interest in minimizing $h(\mu)$ we derive a general dual representation based on optimal transport theory. We have the following theorem (we assume that each \mathcal{X}_i is discrete to avoid any pathological issues):

Theorem 3 (Dual representation) *Assume each \mathcal{X}_i is finite, $i = 1, \dots, n$. We have for all $\mu \in \mathcal{P}^{\otimes}(\mathcal{X})$:*

$$h(\mu) = \sup_v \sum_{i=1}^n \int_{\mathcal{X}_i} v_i(x_i) d\mu_i(x_i) \text{ such that } \forall x \in \mathcal{X}, \sum_{i=1}^n v_i(x_i) \leq H(x_1, \dots, x_n). \quad (6)$$

We denote by $\mathcal{V}(H)$ the set of such potentials v .

Proof This is a usual result in the duality of optimal transport problems [29], which we can prove directly here in the discrete case. Indeed, we have by strong Lagrangian duality (which applies because Slater condition is satisfied):

$$\begin{aligned} h(\mu) &= \inf_{\gamma \in \mathbb{R}_+^{\mathcal{X}}} \sum_{x \in \mathcal{X}} \gamma(x) H(x) \text{ such that } \forall i \in \{1, \dots, n\}, \sum_{x_j \in \mathcal{X}_j, j \neq i} \gamma(x) = \mu_i(x_i) \\ &= \inf_{\gamma \in \mathbb{R}_+^{\mathcal{X}}} \sum_{x \in \mathcal{X}} \gamma(x) H(x) + \sum_{i=1}^n \sum_{x_i \in \mathcal{X}_i} \sup_{v_i \in \mathbb{R}^{x_i}} v_i(x_i) \left[\mu_i(x_i) - \sum_{x_j \in \mathcal{X}_j, j \neq i} \gamma(x) \right] \\ &= \sup_{v_i \in \mathbb{R}^{x_i}, i \in \{1, \dots, n\}} \sum_{i=1}^n \sum_{x_i \in \mathcal{X}_i} v_i(x_i) \mu_i(x_i) \text{ such that } \forall x \in \mathcal{X}, \sum_{i=1}^n v_i(x_i) \leq H(x). \end{aligned}$$

Note that the theory of optimal transport leads to extra constraints on an optimal solution v [25]. \blacksquare

While the duality result above resembles the one for submodular set-functions, it is not exactly equal to it when all \mathcal{X}_i have two elements. For discrete sets, we will make an explicit link in Section 3. We first state that the problem of minimizing $H(x)$ over $x \in \mathcal{X}$, has the following dual

$$\inf_{x \in \mathcal{X}} H(x) = \inf_{\mu \in \mathcal{P}^{\otimes}(\mathcal{X})} h(\mu) = \sup_{V \in \mathcal{V}(H)} \sum_{i=1}^n \inf_{x_i \in \mathcal{X}_i} v_i(x_i). \quad (7)$$

It allows to provide certificates of optimality when minimizing H .

2.5 Separable optimization

In the set-function situation, minimizing the Lovász extension plus a separable convex function has appeared useful in several scenarios [4, Section 8], in particular because a single of such problems leads to a solution to a continuously parameterized set of submodular minimization problems on \mathcal{X} . In our general formulation, the separable convex functions that can be combined are themselves defined through a submodular function and optimal transport.

We thus consider n submodular functions $c_i : \mathcal{X}_i \times [0, 1] \rightarrow \mathbb{R}$, $i = 1, \dots, n$, and define the following function from $\mathcal{P}(\mathcal{X}_i)$ to \mathbb{R} :

$$\varphi_i(\mu_i) = \int_0^1 c_i(F_{\mu_i}^{-1}(t), t) dt. \quad (8)$$

Because c_i is submodular, we have a formulation of φ_i as an optimal transport problem between the measure μ_i on \mathcal{X}_i and the uniform distribution $U[0, 1]$ on $[0, 1]$, as

$$\varphi_i(\mu_i) = \inf_{\gamma_i \in \Pi(\mu_i, U[0,1])} \int_{\mathcal{X}_i \times [0,1]} c_i(x_i, t_i) d\gamma_i(x_i, t_i).$$

It is thus always convex (as as the minimum of a jointly convex problem in γ_i and μ_i). While such functions φ_i may look peculiar, as shown at the end of this section, they include all separable convex functions of F_{μ_i} , that is sums (or integrals) of convex functions of $F_{\mu_i}(x_i) \in \mathbb{R}$.

We now consider the following problem:

$$\inf_{\mu \in \mathcal{P}^{\otimes}(\mathcal{X})} h(\mu) + \sum_{i=1}^n \varphi_i(\mu_i), \quad (9)$$

which is an optimization problem with h , with additional separable transport costs $\varphi_i(\mu_i)$. We may derive a dual problem using the representation from Eq. (6):

$$\begin{aligned} \inf_{\mu \in \mathcal{P}^{\otimes}(\mathcal{X})} h(\mu) + \sum_{i=1}^n \varphi_i(\mu_i) &= \inf_{\mu \in \mathcal{P}^{\otimes}(\mathcal{X})} \sup_{V \in \mathcal{V}(H)} \sum_{i=1}^n \left\{ \int_{\mathcal{X}_i} v_i(x_i) d\mu_i(x_i) + \varphi_i(\mu_i) \right\} \\ &= \sup_{V \in \mathcal{V}(H)} \sum_{i=1}^n \inf_{\mu_i \in \mathcal{P}_i(\mathcal{X}_i)} \left\{ \int_{\mathcal{X}_i} v_i(x_i) d\mu_i(x_i) + \varphi_i(\mu_i) \right\}. \end{aligned} \quad (10)$$

The equation above provides a dual problem to Eq. (9). We may also consider a family of submodular minimization problems, parameterized by $t \in [0, 1]$:

$$\min_{x \in \mathcal{X}} H(x) + \sum_{i=1}^n c_i(x_i, t), \quad (11)$$

with their duals defined from Eq. (7). Note that the two dual problems are defined on the same set $\mathcal{V}(H)$. We can now prove the following theorem relating the two optimization problems.

Theorem 4 (Separable optimization) *Assume H is submodular and all c_i , $i = 1, \dots, n$ are strictly submodular. Then:*

- (a) *If x and x' are minimizers of Eq. (11) for $t > t'$, then $x \leq x'$, i.e., the solutions of Eq. (11) are non-increasing in $t \in [0, 1]$.*

- (b) Given a primal candidate $\mu \in \mathcal{P}^\otimes(\mathcal{X})$ and a dual candidate $v \in \mathcal{V}(H)$, then the duality gap for the problem in Eq. (9) is the integral from $t = 0$ to $t = 1$ of the gaps for the problem in Eq. (11) for the same dual candidate and the primal candidate $(F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)) \in \mathcal{X}$.
- (c) Given the unique solution μ of Eq. (9), for all $t \in [0, 1]$, $(F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)) \in \mathcal{X}$ is a solution of Eq. (11).
- (d) Given solutions $x^t \in \mathcal{X}$ for all problems in Eq. (11), we may define μ through $F_{\mu_i}(x_i) = \sup \{t \in [0, 1], x_i^t \geq x_i\}$, for all i and x_i , so that μ is the optimal solution of Eq. (9).

Proof The first statement is a direct consequence of the strict submodularity of c_i [13]. The main idea is that when we go from t to $t' < t$, then the function difference, i.e., $x_i \mapsto c_i(x_i, t') - c_i(x_i, t)$ is strictly increasing, hence the minimizer has to decrease.

For the second statement, we may first re-write the cost function in Eq. (9) as an integral in t , that is, for any $\mu \in \mathcal{P}^\otimes(\mathcal{X})$:

$$h(\mu) + \sum_{i=1}^n \varphi_i(\mu_i) = \int_0^1 \left\{ H[F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)] + \sum_{i=1}^n c_i[F_{\mu_i}^{-1}(t), t] \right\} dt.$$

The gap for a single submodular minimization problem in Eq. (11) is, for a primal candidate $x \in \mathcal{X}$ and a dual candidate $V \in \mathcal{V}(H)$:

$$H(x) + \sum_{i=1}^n c_i[x_i, F_{\mu_i}^{-1}(t)] - \sum_{i=1}^n \min_{y_i \in \mathcal{X}_i} \left\{ v_i(y_i) + c_i[y_i, t] \right\},$$

and its integral with respect to $t \in [0, 1]$ for $x_i = F_{\mu_i}^{-1}(t)$, for all $i \in \{1, \dots, n\}$, is equal to

$$h(\mu) + \sum_{i=1}^n \varphi_i(\mu_i) - \sum_{i=1}^n \int_0^1 \min_{y_i \in \mathcal{X}_i} \left\{ v_i(y_i) + c_i[y_i, t] \right\} dt.$$

Finally, we have, by using the formulation of $\varphi_i(\mu_i)$ through optimal transport, an expression of the elements appearing in the dual problem in Eq. (10):

$$\begin{aligned} & \inf_{\mu_i \in \mathcal{P}_i(\mathcal{X}_i)} \left\{ \int_{\mathcal{X}_i} v_i(x_i) d\mu_i(x_i) + \varphi_i(\mu_i) \right\} \\ &= \inf_{\mu_i \in \mathcal{P}_i(\mathcal{X}_i)} \inf_{\gamma_i \in \Pi(\mu_i, U[0,1])} \int_{\mathcal{X}_i \times [0,1]} [v_i(x_i) + c_i(x_i, t_i)] d\gamma_i(x_i, t_i) \\ &= \int_0^1 \inf_{y_i \in \mathcal{X}_i} [v_i(y_i) + c_i(y_i, t_i)] dt_i, \end{aligned}$$

because we may for any t_i choose the conditional distribution of x_i given t_i equal to a Dirac at the minimizer y_i of $v_i(y_i) + c_i(y_i, t_i)$. This implies (b) and (c). The statement (d) is proved exactly like for set-functions [4, Prop. 8.3]. ■

Which functions may be represented as submodular optimal transport problems? The formulation in Eq. (8) may look rather restrictive. In fact, it includes all functions of the form

$$\varphi_i(\mu_i) = \int_{\mathcal{X}_i} d\psi(x_i)[F_{\mu_i}(x_i)],$$

where ψ is a measure on \mathcal{X}_i with values convex functions from $[0, 1]$ to \mathbb{R} . When \mathcal{X}_i is discrete, this corresponds to $\varphi_i(\mu_i) = \sum_{x_i \in \mathcal{X}_i} \psi(x_i)[F_{\mu_i}(x_i)]$, where each $\psi(x_i)$ is a convex function from $[0, 1]$ to \mathbb{R} . Indeed, we then have:

$$\begin{aligned} \varphi_i(\mu_i) &= \sum_{x_i \in \mathcal{X}_i} \psi(x_i)(0) + \sum_{x_i \in \mathcal{X}_i} \int_0^1 1_{F_{\mu_i}(x_i) \geq t} \frac{d\psi(x_i)}{dt}(t) dt \\ &= \sum_{x_i \in \mathcal{X}_i} \psi(x_i)(0) + \int_0^1 \sum_{x_i \in \mathcal{X}_i, x_i \leq F_{\mu_i}^{-1}(t)} \frac{d\psi(x_i)}{dt}(t) dt, \end{aligned}$$


and thus, up to constants, we may define $c_i(y_i, t)$ as $\sum_{x_i \in \mathcal{X}_i, x_i \leq y_i} \frac{d\psi(x_i)}{dt}(t)$, for which we have: $\varphi_i(\mu_i) = \int_0^1 c_i(F_{\mu_i}^{-1}(t), t) dt$. In the general continuous case, we could define $c_i(y_i, t)$ similarly as the derivative with respect to t of the mass under ψ of the set $\{t_i \in \mathcal{X}_i, t_i \leq y_i\}$.

3 Discrete sets

In this section, we consider only finite sets for all $i \in \{1, \dots, n\}$, i.e., $\mathcal{X}_i = \{0, \dots, k_i - 1\}$. For this important subcase, we will extend many of the notions related to submodular set-functions, such as the base polytope and the greedy algorithm to compute its support function. This requires to extend the domain where we compute our extensions.

This will be done using a specific representation of the measures $\mu_i \in \mathcal{P}(\mathcal{X}_i)$. Indeed, we may represent μ_i through its cumulative distribution function $\rho_i(x_i) = \mu_i(x_i) + \dots + \mu_i(k_i) = F_{\mu_i}(x_i)$, for $x_i \in \{1, \dots, k_i - 1\}$. Because the measure μ_i has unit total mass, $\rho_i(0)$ is always equal to 0 and can be left out. The only constraint on ρ_i is that it has non-increasing components and that all of them belong to $[0, 1]$. We denote by $[0, 1]_{\downarrow}^{k_i-1}$ this set, which is in bijection with $\mathcal{P}(\{0, \dots, k_i - 1\})$.

We are going to consider the set of non-increasing vectors $\mathbb{R}_{\downarrow}^{k_i-1}$ (without the constraint that they are between 0 and 1). For any such ρ_i with $k_i - 1$ non-increasing components, the set of real numbers is divided into k_i parts, as shown below.



This creates a map $\theta(\rho_i, \cdot) : \mathbb{R} \rightarrow \{0, \dots, k_i - 1\}$ such that $\theta(\rho_i, t) = k_i - 1$ for $t < \rho_i(k_i - 1)$, $\theta(\rho_i, t) = x_i$ if $t \in (\rho_i(x_i + 1), \rho_i(x_i))$, for $x_i \in \{1, \dots, k_i - 2\}$, and $\theta(\rho_i, t) = 0$ for $t > \rho_i(1)$. What happens at the boundary points is arbitrary.

For example, for $k_i = 2$ (and $\mathcal{X}_i = \{0, 1\}$), then we simply have $\rho_i \in \mathbb{R}$ and $\theta(\rho_i, t) = 1$ for $t < \rho_i$, and 0 if $t > \rho_i$. We can now give an expression of $h(\mu)$ as a function of $\rho \in \prod_{i=1}^n [0, 1]_{\downarrow}^{k_i-1}$, which will be extended to all non-increasing vectors (not constrained to be between 0 and 1). This will then allow us to define extensions of base polytopes.

3.1 Extended extension

We first start by a simple lemma providing an expression of $h(\mu)$ as a function of ρ —note the similarity with the Lovász extension for set-functions.

Lemma 1 (Extension as a function of ρ) If $\mu \in \mathcal{P}^{\otimes}(\mathcal{X})$ and $\rho \in \prod_{i=1}^n [0, 1]_{\downarrow}^{k_i-1}$ are linked through $\rho_i(x_i) = F_{\mu_i}(x_i)$, define $h_{\downarrow}(\rho)$ as

$$h_{\downarrow}(\rho) = \int_0^1 H(\theta(\rho_1, t), \dots, \theta(\rho_n, t)) dt. \quad (12)$$

Then, $h(\mu) = h_{\downarrow}(\rho)$.

Proof This is simply a re-writing of the definition of θ , as for almost all t , and all i , we have $\theta(\rho_i, t) = F_{\mu_i}^{-1}(t)$. \blacksquare

We can give an alternative formulation for $h_{\downarrow}(\rho)$ for $\rho \in \prod_{i=1}^n [0, 1]_{\downarrow}^{k_i-1}$, as (with $\max\{\rho\}$ the maximum value of all $\rho_i(x_i)$, and similarly for $\min\{\rho\}$):

$$\begin{aligned} h_{\downarrow}(\rho) &= \left(\int_0^{\min\{\rho\}} + \int_{\min\{\rho\}}^{\max\{\rho\}} + \int_{\max\{\rho\}}^1 \right) H(\theta(\rho_1, t), \dots, \theta(\rho_n, t)) dt \\ &= \int_{\min\{\rho\}}^{\max\{\rho\}} H(\theta(\rho_1, t), \dots, \theta(\rho_n, t)) dt + \min\{\rho\} H(k_1 - 1, \dots, k_n - 1) + (1 - \max\{\rho\}) H(0). \end{aligned} \quad (13)$$

The expression in Eq. (13) may be used as a definition of $h_{\downarrow}(\rho)$ for all $\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}$. With this definition, the function $\rho \mapsto h_{\downarrow}(\rho) - H(0)$ is piecewise linear.

We now provide a simple algorithm to compute $h_{\downarrow}(\rho)$ in Eq. (13), that extends the greedy algorithm for submodular functions. We thus assume that we are given $\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}$, and we compute $h_{\downarrow}(\rho)$ without integration.

We first order all $r = \sum_{i=1}^n k_i - n$ values of $\rho_i(x_i)$ for all $x_i \in \{1, \dots, k_i - 1\}$, in decreasing order, breaking ties randomly, except to ensure that all values for $\rho_i(x_i)$ for a given i are in the correct order (such ties may occur when one associated $\mu_i(x_i)$ is equal to zero).

We assume that the s -th value is equal to $t(s)$ and corresponds to $\rho_{i(s)}(j(s))$. We have $t(1) = \max\{\rho\}$ and $t(r) = \min\{\rho\}$. For $s \in \{1, \dots, r-1\}$, we define the vector $y(s) \in \mathcal{X}$ that will be the value of $\theta(\rho, t)$ on the open interval (potentially empty) $(t(s+1), t(s))$; note that what happens at break points is still irrelevant. By convention, we define $y(0) = (0, \dots, 0)$ and $y(r) = (k_1 - 1, \dots, k_n - 1)$. We then go from $y(s-1)$ to $y(s)$ by increasing the $i(s)$ -th component by one. Note that we always have $y(s)_{i(s)} = j(s)$.

We thus get, by cutting the integral on $[\min\{\rho\}, \max\{\rho\}] = [t(r), t(1)]$ into pieces:

$$\begin{aligned} h_{\downarrow}(\rho) &= \min\{\rho\} H(k_1 - 1, \dots, k_n - 1) + (1 - \max\{\rho\}) H(0) + \sum_{s=1}^{r-1} \int_{t(s+1)}^{t(s)} H[y(s)] dt \\ &= H(0) + t(r) H[y(r)] - t(1) H[y(0)] + \sum_{s=1}^{r-1} [t(s) - t(s+1)] \cdot H[y(s)] \\ &= H(0) + \sum_{s=1}^r t(s) \left(H[y(s)] - H[y(s-1)] \right). \end{aligned}$$

Since we have ordered all $r = \sum_{i=1}^n k_i - n$ values of $t(s) = \rho_{i(s)}(j(s))$, each $\rho_i(x_i)$ appears exactly once in the sum above. Therefore, $h_{\downarrow}(\rho)$ is of the form $h_{\downarrow}(\rho) = H(0) + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \rho_i(x_i) w_i(x_i)$, where $w_i(x_j)$ is a difference of two function values. Moreover, we have $\sum_{i=1}^n \sum_{x_i=1}^{k_i-1} w_i(x_i) = H(k_1 - 1, \dots, k_n - 1) - H(0)$. We refer to this w as the output of the greedy algorithm (associated with a specific ordering of the values of ρ).

Note that for $k_i = 2$ for all $i \in \{1, \dots, n\}$, then we exactly obtain the greedy algorithm for submodular set-functions. The complexity is $O(r \log r)$ for sorting and r function evaluations, with $r = \sum_{i=1}^n k_i - n$.

3.2 Base polyhedron

We may now provide an extension of the base polyhedron which is usually defined for submodular set-functions. As opposed to set-functions, there are two natural polyhedra $\mathcal{W}(H)$ and $\mathcal{B}(H)$, one defined by linear inequalities, one defined as the convex hull of the outputs of the greedy algorithm. They are equal for set-functions, but not in the general case. The key difference is the monotonicity constraint on each ρ_i , which is only active when $k_i > 2$.

We consider the set $\mathcal{W}(H) \subset \prod_{i=1}^n \mathbb{R}^{k_i-1}$ defined through the inequalities:

$$\begin{aligned} \forall (x_1, \dots, x_n) \in \mathcal{X}, \quad & \sum_{i=1}^n \sum_{y_i=1}^{x_i} w_i(y_i) \leq H(x_1, \dots, x_n) - H(0) \\ & \sum_{i=1}^n \sum_{y_i=1}^{k_i-1} w_i(y_i) = H(k_1-1, \dots, k_n-1) - H(0). \end{aligned}$$

Note that when some $x_i = 0$, then the sum $\sum_{y_i=1}^{x_i} w_i(y_i)$ is equal to zero by convention. When $k_i = 2$ for all $i \in \{1, \dots, n\}$, we obtain exactly the usual base polyhedron (which happens to be bounded). The following proposition shows that it behaves in a similar way (except that it is not bounded).

Proposition 1 (Support function for $\mathcal{W}(H)$) *Assume H is submodular. Then, for any $\rho \in \prod_{i=1}^n \mathbb{R}^{k_i-1}$,*

$$\sup_{w \in \mathcal{W}(H)} \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} w_i(x_i) \rho_i(x_i)$$

is equal to $+\infty$ if $\rho \notin \prod_{i=1}^n \mathbb{R}_+^{k_i-1}$, and to $h_{\downarrow}(\rho) - H(0)$ otherwise, with an optimal $w \in \mathcal{W}(H)$ obtained from the greedy algorithm.

Proof In this proof, we are going to use a reparameterization of w as follows; we define a vector $v_i \in \mathbb{R}^{\{1, \dots, k_i-1\}}$ such that $v_i(x_i) = \sum_{y_i=1}^{x_i} w_i(y_i)$ (this is a bijection from w_i to v_i). We then have the constraint that $\sum_{i=1}^n v_i(x_i) \leq H(x_1, \dots, x_n) - H(0)$, for all $x \in \mathcal{X}$. The extra constraint is that $\sum_{i=1}^n v_i(k_i-1) = H(k_1-1, \dots, k_n-1) - H(0)$. Moreover, we have, with $\mu_i(x_i) = \rho_i(x_i) - \rho_i(x_i+1)$ for $x_i \in \{1, \dots, k_i-2\}$ and $\mu_i(k_i-1) = \rho_i(k_i-1)$, an expression for the linear function we aim to maximize:

$$\begin{aligned} \sum_{x_i=1}^{k_i-1} w_i(x_i) \rho_i(x_i) &= \sum_{x_i=2}^{k_i-1} [v_i(x_i) - v_i(x_i-1)] \rho_i(x_i) + v_i(1) \rho_i(1) \\ &= \sum_{x_i=1}^{k_i-2} v_i(x_i) [\rho_i(x_i) - \rho_i(x_i+1)] + v_i(k_i-1) \rho_i(k_i-1) = \sum_{x_i=1}^{k_i-1} v_i(x_i) \mu_i(x_i). \end{aligned}$$

We first show that if one ρ_i is not non-decreasing, then the supremum is equal to infinity. In such a case, then there exists $x_i \in \{1, \dots, k_i-2\}$ such that $\mu_i(x_i) > 0$. We may then let the corresponding $v_i(x_i)$ tend to $-\infty$.

We now assume that each ρ_i is non-decreasing. We are going to follow the same proof than for set-functions. First, given the piecewise-linearity of $h_\downarrow - H(0)$ and the fact that $h_\downarrow(\rho + C) = h_\downarrow(\rho) + C[H(k_1 - 1, \dots, k_n - 1) - H(0)]$, we only need to show the result for $\rho_i(x_i) \in [0, 1]$ for all i , and $x_i \in \mathcal{X}_i$. Each vector ρ_i is then uniquely associated to a probability measure μ_i on $\mathcal{X}_i = \{0, \dots, k_i - 1\}$, and, from Lemma 1, $h_\downarrow(\rho) = h(\mu)$.

Using the parameterization in terms of v and μ , we can now consider the Lagrangian

$$\mathcal{L}(v, \gamma) = \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} v_i(x_i) \mu_i(x_i) + \sum_{x \in \mathcal{X}} \gamma(x) \left(H(x) - H(0) - \sum_{i=1}^n v_i(x_i) \delta(x_i > 0) \right).$$

The dual problem becomes the one of minimizing $\sum_{x \in \mathcal{X}} \gamma(x) H(x)$ such that

$$\begin{aligned} \forall i \in \{1, \dots, n\}, \forall x_i \in \{1, \dots, k_i - 1\}, \quad & \gamma_i(x_i) = \mu_i(x_i) \\ \forall x \in \mathcal{X} \setminus \{(k_1 - 1, \dots, k_n - 1)\}, \quad & \gamma(x) \geq 0, \end{aligned}$$

where γ_i is the marginal of γ on the i -th variable.

We now exhibit primal/dual pairs with equal objective values. For the dual variable γ , we consider the solution of the usual optimal transport problem (which happens to satisfy extra constraints, that is nonnegativity also for $x = 0$), and the dual value is exactly the extension $h(\mu)$. For the primal variable, we consider the w parameterization (which is in bijection with v). From the greedy algorithm described in Section 3.1. The vector w satisfies the sum constraint $\sum_{i=1}^n \sum_{y_i=1}^{k_i-1} w_i(y_i) = H(k_1 - 1, \dots, k_n - 1) - H(0)$. We simply need to show that for all $x \in \mathcal{X}$, $\sum_{i=1}^n \sum_{z_i=1}^{x_i} w_i(z_i) \leq H(x) - H(0)$. We have, using the notation of the greedy algorithm from Section 3.1:

$$\begin{aligned} \sum_{i=1}^n \sum_{z_i=1}^{x_i} w_i(z_i) &= \sum_{s=1}^r (H[y(s)] - H[y(s-1)]) \delta(y(s)_{i(s)} \leq x_{i(s)}) \text{ by definition of } i(s) \text{ and } y(s), \\ &= \sum_{a=1}^n \sum_{s, i(s)=a} (H[y(s)] - H[y(s-1)]) \delta(y(s)_a \leq x_a) \text{ by splitting the values of } i(s), \\ &= \sum_{a=1}^n \sum_{s, i(s)=a} (H[y(s-1) + e_a] - H[y(s-1)]) \delta(y(s)_a \leq x_a) \\ &\quad \text{because we go from } y(s-1) \text{ to } y(s) \text{ by incrementing the component } i(s) = a, \\ &\leq \sum_{a=1}^n \sum_{s, i(s)=a} (H[\min\{y(s-1) + e_a, x\}] - H[\min\{y(s-1), x\}]) \delta(y(s-1)_a + 1 \leq x_a) \\ &\quad \text{by submodularity,} \\ &= \sum_{a=1}^n \sum_{s, i(s)=a} (H[\min\{y(s-1) + e_a, x\}] - H[\min\{y(s-1), x\}]) \\ &\quad \text{because the difference in values of } H \text{ is equal to zero for } y(s-1)_a + 1 > x_a, \\ &= \sum_{s=1}^r (H[\min\{y(s), x\}] - H[\min\{y(s-1), x\}]) \\ &= H[\min\{x, y(r)\}] - H[\min\{x, y(0)\}] = H(x) - H(0). \end{aligned}$$

Thus, w is feasible. By construction the primal value is equal to $h(\mu)$. We thus have a primal/dual optimal pair and the result is proved. \blacksquare

Given the representation of $h_{\downarrow}(\rho) - H(0)$ as a maximum of linear functions (with an infinite value if some ρ_i does not have non-increasing components), the convex problem of minimizing $H(x)$, which is equivalent to minimizing $h_{\downarrow}(\rho) - H(0)$ with respect to $\rho \in \prod_{i=1}^n [0, 1]_{\downarrow}^{k_i-1}$ has the following dual problem:

$$\max_{w \in \mathcal{W}(H)} \sum_{i=1}^n \min_{\rho_i \in [0, 1]_{\downarrow}^{k_i-1}} \sum_{x_i=1}^{k_i-1} w_i(x_i) \rho_i(x_i) = \max_{w \in \mathcal{W}(H)} \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \min\{w_i(x_i), 0\}. \quad (14)$$

While $\mathcal{W}(H)$ share some properties with the base polytope of a submodular set-function, it is not bounded in general and is not the convex hull of all outputs of the greedy algorithm from Section 3.1.

We now define the base polytope $\mathcal{B}(H)$ as the convex hull of all outputs of the greedy algorithm, when going over all allowed orderings of $\sum_{i=1}^n k_i - n$ elements of ρ that respect the ordering of the individual ρ_i 's. In the submodular set-function case, we have $\mathcal{B}(H) = \mathcal{W}(H)$. However, in the general case we only have an inclusion.

Proposition 2 (Properties of the base polytope $\mathcal{B}(H)$) *Assume H is submodular. Then:*

- (a) $\mathcal{B}(H) \subset \mathcal{W}(H)$.
- (b) For any $\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}$, $\max_{w \in \mathcal{B}(H)} \langle w, \rho \rangle = \max_{w \in \mathcal{W}(H)} \langle w, \rho \rangle = h_{\downarrow}(\rho) - H(0)$, with a joint maximizer obtained as the output of the greedy algorithm.
- (c) $\mathcal{W}(H) = \mathcal{B}(H) + \left\{ w \in \prod_{i=1}^n \mathbb{R}^{k_i-1}, \forall i \in \{1, \dots, n\}, \forall x_i \in \{1, \dots, k_i - 1\}, \sum_{y_i=1}^{x_i} w_i(y_i) \leq 0 \right\}$.

Proof In the statements above, $\langle w, \rho \rangle$ stands for $\sum_{i=1}^n \sum_{x_i=1}^{k_i-1} w_i(x_i) \rho_i(x_i)$. The statements (a) and (b) were shown in the proof of the previous proposition. Statement (c) is a simple consequence of the fact that the polar cone to $\prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}$ is what needs to be added to $\mathcal{B}(H)$ to get to $\mathcal{W}(H)$. ■

The key difference between $\mathcal{W}(H)$ and $\mathcal{B}(H)$ is that the support function of $\mathcal{B}(H)$ does not include the constraint that the argument should be composed of non-increasing vectors. However, $\mathcal{B}(H)$ is a polytope (i.e., bounded), while $\mathcal{W}(H)$ is not.

To obtain dual problems, we may either choose $\mathcal{W}(H)$ or $\mathcal{B}(H)$ (and then take into account the monotonicity constraints explicitly). For our algorithms in Section 4, we will consider $\mathcal{B}(H)$, while in the section below, we will use $\mathcal{W}(H)$ for the analysis and comparison of several optimization problems. The dual using $\mathcal{W}(H)$ is given in Eq. (14). When using $\mathcal{B}(H)$, the dual problem of minimizing $h_{\downarrow}(\rho) - H(0)$ with respect to $\rho \in \prod_{i=1}^n [0, 1]_{\downarrow}^{k_i-1}$ has the following dual problem:

$$\max_{w \in \mathcal{B}(H)} \sum_{i=1}^n \min_{\rho_i \in [0, 1]_{\downarrow}^{k_i-1}} \sum_{x_i=1}^{k_i-1} w_i(x_i) \rho_i(x_i) = \max_{w \in \mathcal{B}(H)} \sum_{i=1}^n \min_{x_i \in \{0, \dots, k_i-1\}} \sum_{y_i=1}^{x_i} w_i(y_i). \quad (15)$$

3.3 Separable optimization

We consider convex functions $\psi_{i x_i} : \rho_i(x_i) \mapsto \psi_{i x_i}[\rho_i(x_i)]$, for $i \in \{1, \dots, n\}$ and $x_i \in \{1, \dots, k_i - 1\}$, which are convex on \mathbb{R} . For simplicity, we follow the same assumptions than in [4, Section 8], that is, they are all strictly convex, continuously differentiable and such that the images of their derivatives goes from $-\infty$ to $+\infty$. Their Fenchel conjugates $\psi_{i x_i}^*$ then have full domains and are differentiable.

We consider the pair of primal/dual problems:

$$\min_{\rho} h_{\downarrow}(\rho) + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \psi_{ix_i}[\rho_i(x_i)] \text{ such that } \rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1} \quad (16)$$

$$\begin{aligned} &= \min_{\rho \in \prod_{i=1}^n \mathbb{R}^{k_i-1}} \max_{w \in \mathcal{W}(H)} \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \left\{ w_i(x_i) \rho_i(x_i) + \psi_{ix_i}[\rho_i(x_i)] \right\} \\ &= \max_{w \in \mathcal{W}(H)} \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} -\psi_{ix_i}^*(-w_i(x_i)). \end{aligned} \quad (17)$$

Note that we have used $\mathcal{W}(H)$ instead of $\mathcal{B}(H)$ to include automatically the constraints of monotonicity of ρ . If using $\mathcal{B}(H)$, we would get a formulation which is more adapted to optimization algorithms (see Section 4).

The following theorem relates this optimization problem to the following family of submodular minimization problems, for $t \in \mathbb{R}$:

$$\min_{x \in \mathcal{X}} H(x) + \sum_{i=1}^n \sum_{y_i=1}^{x_i} \psi'_{iy_i}(t). \quad (18)$$

This is a discrete version of Theorem 4, which directly extends the corresponding results from submodular set-functions.

Theorem 5 (Separable optimization – discrete domains) *Assume H is submodular. Then:*

- (a) *If x and x' are minimizers of Eq. (18) for $t > t'$, then $x \leq x'$, i.e., the solutions of Eq. (18) are non-increasing in $t \in \mathbb{R}$.*
- (b) *Given a primal candidate $\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}$ and a dual candidate $w \in \mathcal{W}(H)$ for Eq. (16), then the duality gap for the problem in Eq. (16) is the integral from $t = 0$ to $t = 1$ of the gaps for the problem in Eq. (18) for the same dual candidate and the primal candidate $\theta(\rho, t) \in \mathcal{X}$.*
- (c) *Given the unique solution μ of Eq. (16), for all $t \in \mathbb{R}$, $\theta(\rho, t) \in \mathcal{X}$ is a solution of Eq. (18).*
- (d) *Given solutions $x^t \in \mathcal{X}$ for all problems in Eq. (18), we may define ρ through $\rho_i(x_i) = \sup \{t \in \mathbb{R}, x_i^t \geq x_i\}$, for all i and x_i , so that ρ is the optimal solution of Eq. (16).*

Proof The first statement is a consequence of submodularity [13]. The main idea is that when we go from t to $t' < t$, then the function difference is increasing by convexity, hence the minimizer has to decrease.

For the second statement, we provide expressions for all elements of the gap, following the proof in [4, Prop. 8.5]. For $M > 0$ large enough, we have

$$h_{\downarrow}(\rho) - H(0) = \int_{-M}^{+M} H[\theta(\rho, t)] dt - MH(k_1-1, \dots, k_n-1) - MH(0).$$

Moreover, the integral of the “non-H-dependent” part of the primal objective in for the submodular minimization problem in Eq. (18) is equal to, for $i \in \{1, \dots, n\}$:

$$\begin{aligned}
& \int_{-M}^{+M} \left\{ \sum_{x_i=1}^{k_i-1} \delta(\theta(\rho_i, t) = x_i) \sum_{y_i=1}^{x_i} \psi'_{iy_i}(t) \right\} dt \\
&= \int_{-M}^{\rho_i(k_i-1)} \sum_{y_i=1}^{k_i-1} \psi'_{iy_i}(t) dt + \sum_{a_i=2}^{k_i-1} \int_{\rho_i(a_i)}^{\rho_i(a_i-1)} \sum_{y_i=1}^{a_i-1} \psi'_{iy_i}(t) dt \\
&= \sum_{y_i=1}^{k_i-1} \left\{ \psi_{iy_i}(\rho_i(k-1)) - \psi_{iy_i}(-M) \right\} + \sum_{a_i=2}^{k_i-1} \sum_{y_i=1}^{a_i-1} \left\{ \psi_{iy_i}(\rho_i(a_i-1)) - \psi_{iy_i}(\rho_i(a_i)) \right\} \\
&= - \sum_{y_i=1}^{k_i-1} \psi_{iy_i}(-M) + \sum_{x_i=1}^{k_i-1} \psi_{ix_i}(\rho_i(x_i)).
\end{aligned}$$

We may now compute for any $i \in \{1, \dots, n\}$ and $x_i \in \{0, \dots, k_i - 1\}$, the necessary pieces for the integral of the dual objective values from Eq. (14):

$$\begin{aligned}
& \int_{-M}^M \min\{w_i(x_i) + \psi'_{ix_i}(t), 0\} dt \\
&= \int_{-M}^{(\psi_{ix_i}^*)'(-w_i(x_i))} (w_i(x_i) + \psi'_{ix_i}(t)) dt \\
&= \psi_{ix_i}((\psi_{ix_i}^*)'(-w_i(x_i))) - \psi_{ix_i}(-M) + w_i(x_i)[M + (\psi_{ix_i}^*)'(-w_i(x_i))] \\
&= (\psi_{iy_i}^*)'(-w_i(x_i))(-w_i(x_i)) - \psi_{iy_i}^*(-w_i(x_i)) - \psi_{iy_i}(-M) + w_i(x_i)[M + (\psi_{ix_i}^*)'(-w_i(x_i))] \\
&= -\psi_{ix_i}^*(-w_i(x_i)) - \psi_{ix_i}(-M) + w_i(x_i)M.
\end{aligned}$$

By putting all pieces together, we obtain the desired result, that is,

$$h_{\downarrow}(\rho) - H(0) + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \psi_{ix_i}^*(-w_i(x_i))$$

is the integral from $-M$ to M of

$$H[\theta(\rho, t)] + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \delta(\theta(\rho_i, t) = x_i) \sum_{y_i=1}^{x_i} \psi'_{iy_i}(t) - H(0) - \sum_{i=1}^n \min\{w_i(x_i) + \psi'_{ix_i}(t), 0\}.$$

The proofs of statements (c) and (d) follow the same argument as for Theorem 4. ■

In the next section, we will use this result for the functions $\varphi_{ix_i}(t) = \frac{1}{2}t^2$, and from thresholding at $t = 0$ we obtain a solution for the minimization of H , thus directly extending the submodular set-function situation.

3.4 Relationship with submodular set-functions and ring families

For finite sets, it is known that one may reformulate the submodular minimization problem in terms of a submodular function with added constraints [5]. Given the submodular function H , defined on $\mathcal{X} = \prod_{i=1}^n \{0, \dots, k_i - 1\}$ we consider a submodular set-function defined on a *ring family* of the set

$$V = \{(i, x_i), i \in \{1, \dots, n\}, x_i \in \{0, \dots, k_i - 1\}\} \subset \{1, \dots, n\} \times \mathbb{R},$$

that is, a family of subsets of V which is invariant by intersection and union. In our context, a member of the ring family is such that if (i, x_i) is in the set, then all (i, y_i) , for $y_i \leq x_i$ are in the set as well.

With our notations in this paper, any element of the ring family is uniquely associated to a vector $z \in \prod_{i=1}^n \{0, 1\}_{\downarrow}^{k_i-1}$ (note that we impose that each vector z_i has non-increasing components). Any such element z is also associated uniquely to an element $x \in \mathcal{X} = \prod_{i=1}^n \{0, \dots, k_i - 1\}$, by taking x_i as the largest $y_i \in \{1, \dots, k_i - 1\}$ such that $z_i(y_i) = 1$, and with value zero, if $z_i = 0$. We can thus define a function $H_{\text{ring}}(z)$ equal to $H(x)$ for this uniquely associated x . Then for any two $z, z' \in \prod_{i=1}^n \{0, 1\}_{\downarrow}^{k_i-1}$, $H_{\text{ring}}(z) + H_{\text{ring}}(z') \geq H_{\text{ring}}(\max\{z, z'\}) + H_{\text{ring}}(\min\{z, z'\})$, i.e., H_{ring} is submodular on the ring family, and minimizing $H(x)$ for $x \in \mathcal{X}$, is equivalent to minimizing H_{ring} on the ring family.

There is then a classical reduction to the minimization of a submodular function on $\prod_{i=1}^n \{0, 1\}^{k_i-1}$ (without monotonicity constraints) [5, 14], which is a regular submodular set-function minimization problem on a set of size $\sum_{i=1}^n k_i - n$. For a certain $\alpha > 0$, define the function

$$H_{\text{ring}}^{\text{ext}}(z) = H_{\text{ring}}(z^{\downarrow}) + \alpha \|z^{\downarrow} - z\|_1,$$

where for $z \in \prod_{i=1}^n \{0, 1\}^{k_i-1}$, z^{\downarrow} denotes the smallest element of the ring family containing z ; in other words $z_i^{\downarrow}(x_i) = 1$ for all x_i such that there exists $y_i \geq x_i$ with $z_i(y_i) = 1$. If α is large enough¹, the function $H_{\text{ring}}^{\text{ext}}$ is submodular; moreover, for any $\alpha > 0$, any minimizer of $H_{\text{ring}}^{\text{ext}}$ belongs to the ring family, and hence leads to a minimizer of H .

The Lovász extension of $H_{\text{ring}}^{\text{ext}}$ happens to be equal to, for $\nu \in \prod_{i=1}^n [0, 1]^{k_i-1}$,

$$h_{\downarrow}(\nu^{\downarrow}) + \alpha \sum_{i=1}^n \sum_{x_i=1}^{k_i-2} (\nu_i(x_{i+1}) - \nu_i(x_i))_+,$$

where ν_i^{\downarrow} is the smallest non-increasing vector greater or equal to ν_i . When α tends to $+\infty$, we recover our convex relaxation from a different point of view. Note that in practice, unless we are in special situations like min-cut/max-flow problems [14] (where we can take $\alpha = +\infty$), this strategy adds an extra-parameter α and does not lead to our new interpretations and convex relaxations.

4 Optimization for discrete sets

In this section, we assume that we are given a submodular function H on $\prod_{i=1}^n \{0, \dots, k_i - 1\}$, which we can query through function values (and usually through the greedy algorithm). We assume that for all $x \in \mathcal{X}$, when $x_i < k_i - 1$, $|H(x + e_i) - H(x)|$ is bounded by B_i .

We present algorithms to minimize H . These can be used in continuous domains by discretizing each \mathcal{X}_i . The key is that the overall complexity remains polynomial in n , and the dependence in k_i is weak enough to easily reach high precisions.

4.1 Optimizing on measures

We have the first equivalent formulations:

$$\min_{\mu \in \prod_{i=1}^n \mathcal{P}(\{0, \dots, k_i - 1\})} h(\mu) \Leftrightarrow \min_{\rho \in \prod_{i=1}^n [0, 1]_{\downarrow}^{k_i-1}} h_{\downarrow}(\rho).$$

¹As shown in [5], we simply need that for all $x \leq y$, $|H(y) - H(x)| \leq \alpha \|x - y\|_1$.

Once we get an approximately optimal solution ρ , we compute $\theta(\rho, t)$ for all $t \in [0, 1]$, and select the one with minimal value (this is a by-product of the greedy algorithm, i.e., with the notation of Section 3.1, $\min_s H[y(s)]$). If we have a dual candidate $w \in \mathcal{B}(H)$, we can use Eq. (15) to obtain a certificate of optimality.

We consider the *projected subgradient method* on ρ . We may compute a subgradient of h_\downarrow in $\mathcal{B}(H)$ by the greedy algorithm, and then use n isotonic regressions to perform the n independent orthogonal projection of $\rho - \gamma h'_\downarrow(\rho)$, using the pool-adjacent-violator algorithm [30]. Each iteration has thus complexity $\left(\sum_{i=1}^n k_i\right) \log\left(\sum_{i=1}^n k_i\right)$, which corresponds to the greedy algorithm (which is here the bottleneck).

In terms of convergence rates, each element $w_i(x_i)$ of a subgradient is in $[-B_i, B_i]$. Moreover, for any two elements ρ, ρ' of $\prod_{i=1}^n [0, 1]_\downarrow^{k_i-1}$, we have $\|\rho_i - \rho'_i\|_\infty \leq 1$. Thus, in ℓ_2 -norm, the diameter of the domain of optimization is less than $\sqrt{\sum_{i=1}^n k_i}$, and the norm of subgradients less than $\sqrt{\sum_{i=1}^n k_i B_i^2}$. Thus, the distance to optimum (measured in function values) after t steps is less than [31]

$$\sqrt{\frac{\left(\sum_{i=1}^n k_i\right)\left(\sum_{i=1}^n k_i B_i^2\right)}{t}}.$$

Thus if all B_i and k_i are equal, we get a complexity of $O(nk \log(nk))$ per iteration and a convergence rate of $O(nkB/\sqrt{t})$. In practice, we choose the Polyak rule for the step-size, since we have candidates for the dual problem; that is, we use $\gamma = \frac{h_\downarrow(\rho) - D}{\|w\|_2^2}$, where D is the best dual value so far, which we can obtain with dual candidates which are the averages of all elements of $\mathcal{B}(H)$ seen so far.

Note that when discretizing a continuous function with k labels, then $B \propto 1/k$ and thus the number of iterations is independent of the number of grid steps.

Once we have a primal candidate $\rho \in \prod_{i=1}^n [0, 1]_\downarrow^{k_i-1}$ and a dual candidate $w \in \mathcal{B}(H)$, we may compute the minimum value of H along the greedy algorithm, as a primal value, and

$$\min_{\rho \in \prod_{i=1}^n [0, 1]_\downarrow^{k_i-1}} \langle w, \rho \rangle = \sum_{i=1}^n \min_{x_i \in \{0, \dots, k_i-1\}} \sum_{y_i=1}^{x_i} w_i(y_i),$$

as the dual value.

Finally, note that the projected subgradient descent directly applies when a noisy oracle for H is available, using the usual arguments for stochastic extensions [31], with a similar convergence rate.

4.2 Smooth extension and Frank-Wolfe techniques

We consider the minimization of $h_\downarrow(\rho) + \frac{1}{2}\|\rho\|^2$, with $\|\rho\|^2 = \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} |\rho_i(x_i)|^2$. Given an approximate ρ , we compute $\theta(\rho, t)$ for all $t \in \mathbb{R}$ (which can be done by ordering all values of $\rho_i(x_i)$, like in the greedy algorithm).

We have:

$$\begin{aligned} \min_{\rho \in \prod_{i=1}^n \mathbb{R}_\downarrow^{k_i-1}} h_\downarrow(\rho) + \frac{1}{2}\|\rho\|^2 &= \min_{\rho \in \prod_{i=1}^n \mathbb{R}_\downarrow^{k_i-1}} \max_{w \in \mathcal{B}(H)} \langle \rho, w \rangle + \frac{1}{2}\|\rho\|^2 \\ &= \max_{w \in \mathcal{B}(H)} \left\{ \min_{\rho \in \prod_{i=1}^n \mathbb{R}_\downarrow^{k_i-1}} \langle \rho, w \rangle + \frac{1}{2}\|\rho\|^2 \right\}. \end{aligned}$$




Figure 2: One-dimensional signals: noisy input obtained by adding Gaussian noise to a noiseless signal (left); denoised signal (right).

We thus need to maximize with respect to w in a compact set of diameter less than $2\sqrt{\sum_{i=1}^n k_i B_i^2}$, a 1-smooth function. Thus, we have that after t steps, the distance to optimum is less than $\frac{2}{t} \sum_{i=1}^n k_i B_i^2$ using Frank-Wolfe techniques, with either line-search or fixed step-sizes [32, 33, 34]. Note that primal candidates may also be obtained with a similar convergence rate; the running-time complexity is the same as for subgradient descent in Section 4.1.

Note that when using a form of line-search, we can use warm-restarts, which could be useful when solving a sequence of related problems, e.g., when discretizing a continuous problem.

5 Experiments

In this section, we consider a basic experiment to illustrate our results. We consider the function defined on $[-1, 1]^n$,

$$H(x) = \frac{1}{2} \sum_{i=1}^n (x_i - z_i)^2 + \lambda \sum_{i=1}^n |x_i|^\alpha + \mu \sum_{i=1}^{n-1} (x_i - x_{i+1})^2.$$

This corresponds to denoising a one-dimensional signal z with the constraint that z is smooth and sparse. See an illustration in Figure 2. The smoothness prior is obtained from the quadratic form (which is submodular and convex), while the sparse prior is only submodular, and not convex. Thus, this is not a convex optimization problem; however, it can be solved globally to arbitrary precision using the algorithms presented in Section 4. Note that the use of a non-convex sparse prior leads to fewer biasing artefacts than the usual ℓ_1 -norm [35].

In Figure 3, we show certified duality gaps for the two algorithms from Section 4 on a discretization with 50 grid elements for each of the $n = 50$ variables. We can see that the Frank-Wolfe-based optimization performs better than projected subgradient descent.

Finally, in Figure 4, we show the estimated values for the vectors ρ , showing that the solution is almost a threshold function for the non-smooth dual problem used in Section 4.1 (left), while it provides more information in the smooth dual case used in Section 4.2 (right).

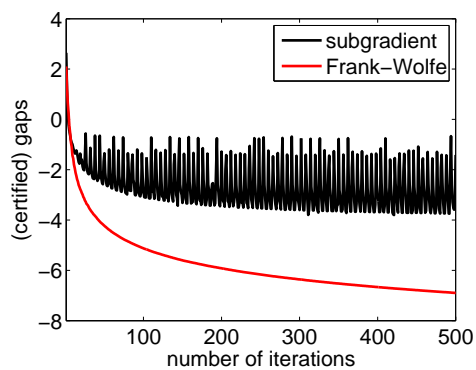


Figure 3: Certified duality gaps: noisy input obtained by adding Gaussian noise to a noiseless signal (left); denoised signal (right).




Figure 4: Optimal values of ρ : non-smooth dual problem (left); smooth dual problem (right). Both axes have been affinely normalized to $[-1, 1]$.

6 Discussion

In this paper, we have shown that a large family of non-convex problems could be solved by a well-defined convex relaxation and efficient algorithms based on simple oracles. This is based on replacing each of the variables x_i by a measure and optimizing over the cumulative distributions of these measures. Our algorithms apply to all submodular functions defined on products of subsets of \mathbb{R}^n , and hence include continuous domains as well as finite domains. Several extensions are worth considering:

- **Relationship with convexity for continuous domains:** for functions defined on a product of sub-intervals, there are two notions of “simple” functions, convex and submodular. These two notions are usually disjoint (see Section 2.1 for examples). It is worth considering these two notions *simultaneously*, as many common objective functions are the sums of a convex function and a submodular function (e.g., the negative log-likelihood of a Gaussian vector, where the covariance matrix is parameterized as a linear combination of positive definite matrices). We can thus consider the minimization of $H(x) + G(x)$, where $H : \mathbb{R}^n \rightarrow \mathbb{R}$ is submodular and $G : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex. We have a natural convex relaxation on our set of measures:

$$\min_{\mu \in \mathcal{P}^{\otimes}(X)} h(\mu) + G\left(\int_{\mathcal{X}} x \prod_{i=1}^n d\mu_i(x_i)\right),$$

which is equivalent to the minimization of the convex function $\tilde{H}(x) + G(x)$ with

$$\tilde{H}(x) = \inf_{\mu \in \mathcal{P}^{\otimes}(X)} h(\mu) \text{ such that } \int_{\mathcal{X}} y \prod_{i=1}^n d\mu_i(y_i) = x.$$

This provides a convex relaxation of $H(x)$, which happens to be the Fenchel bi-dual (i.e., the convex envelope) of H , which is computable since one can minimize $H(x)$ plus linear functions (which are separable). Indeed, we have, by convex duality:

$$\begin{aligned} \tilde{H}(x) &= \inf_{\mu \in \mathcal{P}^{\otimes}(X)} h(\mu) \text{ such that } \int_{\mathcal{X}} y \prod_{i=1}^n d\mu_i(y_i) = x \\ &= \sup_{z \in \mathbb{R}^n} \inf_{\mu \in \mathcal{P}^{\otimes}(X)} h(\mu) + z^\top x - \int_{\mathcal{X}} z^\top y \prod_{i=1}^n d\mu_i(y_i) \\ &= \sup_{z \in \mathbb{R}^n} \inf_{y \in \mathcal{X}} H(y) + z^\top x - y^\top z = H^{**}(x). \end{aligned}$$

The main added benefit of submodularity is that the convex envelope can be computed when H is submodular, whereas typically, it is not possible.

- **Submodular relaxations:** it is possible to write most functions as the difference of two submodular functions H and G , leading to an exact reformulation in terms of minimizing $h(\mu) - g(\mu)$ with respect to the product of measures μ . This problem is however non-convex anymore in general, and we could use majorization-minimization procedures [36]. Alternatively, if a function is a sum of simple functions, we may consider “submodular relaxations” of each of the component (a situation comparable with replacing functions by their convex envelopes). Like in the set-function case, a notion of submodular envelope similar to the convex envelope is not available. However, for any function H , one can always define a convex extension $\tilde{h}(\mu)$ through an optimal transport problem as $\tilde{h}(\mu) = \inf_{\gamma \in \Pi(\mu)} \int_{\mathcal{X}} H(x) d\gamma(x)$. For functions of two variables, this can often be computed in closed form, and by taking the sum of these relaxations, we exactly get the usual linear programming relaxation (see [17] and references therein).

- **Submodular function maximization:** While this paper has focused primarily on minimization, it is worth exploring if algorithms for the maximization of submodular set-functions can be extended to the general case [8, 9].
- **Divide-and-conquer algorithm:** For submodular set-functions, the separable optimization problem defined in Section 2.5 can be exactly solved by a sequence of at most n submodular optimization problem by a divide-and-conquer procedure [37]. It turns out that a similar procedure extends to general submodular functions on discrete domains (see Appendix B).
- **Active-set methods:** For submodular set-functions, active-set techniques such as the minimum-norm-point algorithm have the potential to find the exact minimizer in finitely many iterations [38, 20]; they are based on the separable optimization problem from Section 2.5. Given that our extension simply adds inequality constraints, such an approach could easily be extended.
- **Adaptive discretization schemes:** Faced with functions defined on continuous domains, currently the only strategy is to discretize the domains; it would be interesting to study adaptive discretization strategies based on duality gap criteria.

References

- [1] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001.
- [2] A. Krause and C. Guestrin. Submodularity and its applications in optimized information gathering. *ACM Transactions on Intelligent Systems and Technology*, 2(4), 2011.
- [3] H. Lin and J. Bilmes. A class of submodular functions for document summarization. In *NAACL/HLT*, 2011.
- [4] F. Bach. *Learning with Submodular Functions: A Convex Optimization Perspective*, volume 6 of *Foundations and Trends in Machine Learning*. NOW, 2013.
- [5] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory, Series B*, 80(2):346–355, 2000.
- [6] S. Iwata, L. Fleischer, and S. Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *Journal of the ACM*, 48(4):761–777, 2001.
- [7] J. B. Orlin. A faster strongly polynomial time algorithm for submodular function minimization. *Mathematical Programming*, 118(2):237–251, 2009.
- [8] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294, 1978.
- [9] U. Feige, V. S. Mirrokni, and J. Vondrák. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- [10] G. Choquet. Theory of capacities. *Annales de l’Institut Fourier*, 5:131–295, 1954.
- [11] L. Lovász. Submodular functions and convexity. *Mathematical programming: The state of the art, Bonn*, pages 235–257, 1982.
- [12] G. G. Lorentz. An inequality for rearrangements. *American Mathematical Monthly*, 60(3):176–179, 1953.

- [13] D. M. Topkis. Minimizing a submodular function on a lattice. *Operations research*, 26(2):305–321, 1978.
- [14] H. Ishikawa. Exact optimization for markov random fields with convex priors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1333–1336, 2003.
- [15] T. Pock, A. Chambolle, D. Cremers, and H. Bischof. A convex relaxation approach for computing minimal partitions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 810–817. IEEE, 2009.
- [16] E. Bae, J. Yuan, X.-C. Tai, and Y. Boykov. A fast continuous max-flow approach to non-convex multi-labeling problems. In *Efficient Algorithms for Global Optimization Methods in Computer Vision*, pages 134–154. Springer, 2014.
- [17] S. Živný, T. Werner, and D. Průša. The power of lp relaxation for map inference. In S. Nowozin, P. V. Gehler, J. Jancsary, and C. H. Lampert, editors, *Advanced Structured Prediction*. MIT Press, 2014.
- [18] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [19] J. M. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer, 2006.
- [20] S. Fujishige. *Submodular Functions and Optimization*. Elsevier, 2005.
- [21] S. Friedland and S. Gaubert. Submodular spectral functions of principal submatrices of a hermitian matrix, extensions and applications. *Linear Algebra and its Applications*, 438(10):3872–3884, 2013.
- [22] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the fortieth annual ACM symposium on Theory of Computing*, pages 67–74. ACM, 2008.
- [23] J. Djolonga and A. Krause. From MAP to marginals: Variational inference in Bayesian submodular models. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [24] J. Djolonga and A. Krause. Scalable variational inference in log-supermodular models. In *International Conference on Machine Learning (ICML)*, 2015.
- [25] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [26] L. S. Shapley. Complements and substitutes in the optimal assignment problem. *Naval Research Logistics Quarterly*, 9(1):45–48, 1962.
- [27] B. Pass. Multi-marginal optimal transport: theory and applications. Technical Report 1406.0026, arXiv, 2014.
- [28] G. Carlier. On a class of multidimensional optimal transportation problems. *Journal of Convex Analysis*, 10(2):517–530, 2003.
- [29] H. G. Kellerer. Duality theorems for marginal problem. *Journal of Probability Theory and Related Fields*, 67(4):399–432, 1984.
- [30] M. J. Best and N. Chakravarti. Active set algorithms for isotonic regression: a unifying framework. *Mathematical Programming*, 47(1):425–439, 1990.

- [31] N. S. Shor. *Nondifferentiable optimization and polynomial problems*, volume 24. Springer Science & Business Media, 2013.
- [32] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
- [33] M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2013.
- [34] F. Bach. Duality between subgradient and conditional gradient methods. *SIAM Journal on Optimization*, 25(1):115–129, 2015.
- [35] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- [36] K. Lange, D. R. Hunter, and I. Yang. Optimization transfer using surrogate objective functions. *Journal of computational and graphical statistics*, 9(1):1–20, 2000.
- [37] H. Groenevelt. Two algorithms for maximizing a separable concave function over a polymatroid feasible region. *European Journal of Operational Research*, 54(2):227–236, 1991.
- [38] P. Wolfe. Finding the nearest point in a polytope. *Mathematical Programming*, 11(1):128–149, 1976.

A Proof of miscellaneous results

A.1 Submodularity of the Lovász extension of a submodular set-function

We consider a submodular set-function $G : \{1, \dots, n\} \rightarrow \mathbb{R}$ and its Lovász extension $g : \mathbb{R}^n \rightarrow \mathbb{R}$. In order to show the submodularity of g , we simply apply the definition and consider $w \in \mathbb{R}^n$ and two distinct basis vectors e_i and e_j with infinitesimal positive displacements a_i and a_j . If (i, j) belongs to two different level sets of w , then, for a_i and a_j small enough, $g(w + a_i e_i) + g(w + a_j e_j) - g(w) - g(w + a_i e_i + a_j e_j)$ is equal to zero. If (i, j) belongs to the same level sets, then, by explicitly computing the quantity for $a_i > a_j$ and $a_i < a_j$, it is non-negative (as a consequence of submodularity).

Note that then, the extension has another expression when $\mathcal{X} = [0, 1]^n$, as, if $H = g$,

$$\begin{aligned}
 h(\mu) &= \int_0^1 g(F_{\mu_1}^{-1}(t), \dots, F_{\mu_n}^{-1}(t)) dt \\
 &= \int_0^1 \int_0^1 G(\{i, F_{\mu_i}^{-1}(t) \geq z\}) dt dz = \int_0^1 \int_0^1 G(\{i, F_{\mu_i}(z) \geq t\}) dt dz \\
 &= \int_0^1 g(F_{\mu_1}(z), \dots, F_{\mu_2}(z)) dz,
 \end{aligned}$$

which provides another proof of submodularity.

A.2 Submodularity of the multi-linear extension of a submodular set-function

We consider a submodular set-function $G : \{1, \dots, n\} \rightarrow \mathbb{R}$ and its multi-linear extension [22], defined as a function $\tilde{g} : [0, 1]^n \rightarrow \mathbb{R}$ with

$$\tilde{g}(x_1, \dots, x_n) = \mathbb{E}_{y_i \sim \text{Bernoulli}(x_i), i \in \{1, \dots, n\}} G(y),$$

where all Bernoulli random variables y_i are independent. In order to show submodularity, we only need to consider the case $n = 2$, for which the function \tilde{g} is quadratic in x and the cross-term is non-positive because of the submodularity of G . See more details in [22]. The extension on a product of measures on $[0, 1]^n$ does not seem to have the same simple interpretation as for the Lovász extension in Appendix A.1.

B Divide-and-conquer algorithm for separable optimization

We consider the optimization problem studied in Section 3.3:

$$\min_{\rho} h_{\downarrow}(\rho) + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \psi_{ix_i}[\rho_i(x_i)] \text{ such that } \rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1},$$

which is dual to the problem of maximizing $\sum_{i=1}^n \sum_{x_i=1}^{k_i-1} -\psi^*(-w_i(x_i))$ such that $w \in \mathcal{W}(H)$. We know that it is equivalent to a sequence of submodular minimization problems of the form:

$$\min_{x \in \mathcal{X}} H(x) + \sum_{i=1}^n \sum_{y_i=1}^{x_i} \psi'_{iy_i}(t).$$

We now show that by solving a sequence of such problems (with added restrictions on the domains of the variables), we recover exactly the solution of the original problem. This algorithm directly extends the one from [37], and we follow the exposition from [4, Section 9.1]. For simplicity, we consider strictly convex function differentiable functions ψ_{ix_i} such that their Fenchel-conjugates are defined on \mathbb{R} . The recursive algorithm is as follows:

- (1) Find the unique global maximizer of $-\sum_{i=1}^n \sum_{x_i=1}^{k_i-1} \psi^*(-w_i(x_i))$ with the constraint that $\sum_{i=1}^n \sum_{x_i=1}^{k_i-1} w_i(x_i) = H(k_1, \dots, k_n) - H(0)$.
- (2) Find an element $y \in \mathcal{X}$ that minimizes $H(y) - \sum_{i=1}^n \sum_{z_i=1}^{y_i} w_i(z_i)$.
- (3) If $H(y) - \sum_{i=1}^n \sum_{z_i=1}^{y_i} w_i(z_i) = H(0)$, then exit (w is optimal).
- (4) Maximize $\sum_{i=1}^n \sum_{x_i=1}^{y_i} -\psi^*(-w_i(x_i))$ over $w \in \mathcal{W}(H_{x \leq y})$, where $H_{x \leq y}$ is the restriction of H to $\{x \leq y\}$, to obtain $w^{\{x \leq y\}}$.
- (5) Maximize $\sum_{i=1}^n \sum_{x_i=y_i+1}^{k_i-1} -\psi^*(-w_i(x_i))$ over $w \in \mathcal{W}(H_{x \geq y})$, where $H_{x \geq y}$ is the restriction of H to $\{x \geq y\}$, to obtain $w^{\{x \geq y\}}$. Note that the index x_i then does actually start from $y_i + 1$.
- (6) Concatenate the two vectors $w^{\{x > y\}}$ and $w^{\{x \leq y\}}$ into w , which is the optimal solution.

The proof of correctness is the same as for set-functions: if the algorithm stops at step (3), then we indeed have the optimal solution because the optimum on a wider set happens to be in $\mathcal{W}(H)$. If not, we first show that w obtained in (6) is in $\mathcal{W}(H)$, which is a consequence of submodularity. The global optimality is then due to the fact that the continuous extension $h_{\downarrow}(\nu)$ decouples into the sum of two terms that depend on disjoint sets.