



Generalized conditional gradient: analysis of convergence and applications

Alain Rakotomamonjy, Rémi Flamary, Nicolas Courty

► **To cite this version:**

Alain Rakotomamonjy, Rémi Flamary, Nicolas Courty. Generalized conditional gradient: analysis of convergence and applications. [Research Report] LITIS; Lagrange; IRISA. 2015.

HAL Id: hal-01217870

<https://hal.archives-ouvertes.fr/hal-01217870>

Submitted on 20 Oct 2015

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.



Distributed under a Creative Commons Attribution 4.0 International License

Generalized conditional gradient: analysis of convergence and applications

Alain Rakotomamonjy
Université de Rouen
France
alain.rakoto@insa-rouen.fr

Rémi Flamary
Lagrange, CNRS, UNS, OCA
France
remi.flamary@unice.fr

Nicolas Courty
Université de Bretagne du Sud, IRISA
France
ncourty@irisa.fr

Abstract

The objectives of this technical report is to provide additional results on the generalized conditional gradient methods introduced by Bredies et al. [BLM05]. Indeed, when the objective function is smooth, we provide a novel certificate of optimality and we show that the algorithm has a linear convergence rate. Applications of this algorithm are also discussed.

1 Generalized conditional gradient

We are interested in the problem of minimizing under constraints a composite function such as

$$\min_{\mathbf{x} \in \mathcal{P}} F(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x}), \quad (1)$$

where both $f(\cdot)$ and $g(\cdot)$ are convex and differentiable functions and \mathcal{P} is compact set of \mathbb{R}^n . One might want to benefit from this composite structure during the optimization procedure. For instance, if we have an efficient solver for optimizing

$$\min_{\mathbf{x} \in \mathcal{P}} \langle \nabla f, \mathbf{x} \rangle + g(\mathbf{x}) \quad (2)$$

it is of prime interest to use this solver in the optimization scheme instead of linearizing the whole objective function as one would do with a conditional gradient algorithm [Ber99].

The resulting approach is defined in Algorithm 1, denoted CGS in the remainder (the S standing for Splitting). Conceptually, this algorithm lies in-between the original optimization problem and the conditional gradient. Indeed, if we do not consider any linearization, the step 3 of the algorithm is equivalent to solving the original problem and one iterate will suffice for convergence. If we use a full linearization as in the conditional gradient approach, step 3 is equivalent to solving a rough approximation of the original problem. By linearizing only a part of the objective function, we optimize a better approximation of that function, as opposed to a full linearization as in the conditional gradient approach. This leads to a provably better certificate of optimality than the one of the conditional gradient algorithm [Jag13]. Also, if an efficient solver of the partially linearized problem is available, this algorithm is of strong interest. This is notably the case in computation of regularized optimal transport problems [CFTR15], or Elastic-net regularization [ZH05], see the application Section for details.

Note that this partial linearization idea has already been introduced by Bredies et al. [BLM05] for solving problem 1. Their theoretical results related to the resulting algorithm apply when $f(\cdot)$ is

Algorithm 1 Generalized Gradient Splitting (CGS)

- 1: Initialize $k = 0$ and $\mathbf{x}^0 \in \mathcal{P}$
- 2: **repeat**
- 3: Solve problem $\mathbf{s}^k = \operatorname{argmin}_{\mathbf{x} \in \mathcal{P}} \langle \nabla f(\mathbf{x}^k), \mathbf{x} \rangle + g(\mathbf{x})$
- 4: Find the optimal step α^k with $\Delta \mathbf{x} = \mathbf{s}^k - \mathbf{x}^k$

$$\alpha^k = \operatorname{argmin}_{0 \leq \alpha \leq 1} f(\mathbf{x}^k + \alpha \Delta \mathbf{x}) + g(\mathbf{x}^k + \alpha \Delta \mathbf{x})$$

or choose α^k so that it satisfies the Armijo rule.

- 5: $\mathbf{x}^{k+1} \leftarrow \mathbf{x}^k + \alpha^k \Delta \mathbf{x}$, set $k \leftarrow k + 1$
 - 6: **until** Convergence
-

differentiable, $g(\cdot)$ convex f and g satisfy some others mild conditions like coercivity. These results state that the generalized conditional gradient algorithm is a descent method and that any limit point of the algorithm is a stationary point of $f + g$.

In what follows we provide some results when f and g are differentiable. Some of these results provide novel insights on the generalized gradient algorithms (relation between optimality and minimizer of the search direction, convergence rate, and optimality certificate; while some are redundant to those proposed by Bredies (convergence).

2 Convergence analysis

Before discussing convergence of the algorithm, we first reformulated its step 3 so as to make its properties more accessible and its convergence analysis more amenable.

The reformulation we propose is

$$\mathbf{s}^k = \operatorname{argmin}_{\mathbf{s} \in \mathcal{P}} \langle \nabla f(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle + g(\mathbf{s}) - g(\mathbf{x}^k) \quad (3)$$

and it is easy to note that the problem in line 3 of Algorithm 1, is equivalent to this one and leads to the same solution.

The above formulation allows us to derive a property that highlights the relation between problems 1 and 3.

Proposition 2.1 \mathbf{x}^* is a minimizer of problem (1) if and only if

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{s} \in \mathcal{P}} \langle \nabla f(\mathbf{x}^*), \mathbf{s} - \mathbf{x}^* \rangle + g(\mathbf{s}) - g(\mathbf{x}^*) \quad (4)$$

Proof The proof relies on optimality conditions of constrained convex optimization problem. Indeed, for a convex and differentiable f and g , \mathbf{x}^* is solution of problem (1) if and only if [BNO03]

$$-\nabla f(\mathbf{x}^*) - \nabla g(\mathbf{x}^*) \in N_{\mathcal{P}}(\mathbf{x}^*) \quad (5)$$

where $N_{\mathcal{P}}(\mathbf{x})$ is the normal cone of \mathcal{P} at \mathbf{x} . In a same way, a minimizer \mathbf{s}^* of problem (3) at \mathbf{x}^k can also be characterized as

$$-\nabla f(\mathbf{x}^k) - \nabla g(\mathbf{s}^*) \in N_{\mathcal{P}}(\mathbf{s}^*) \quad (6)$$

Now suppose that \mathbf{x}^* is a minimizer of problem (1), it is easy to see that if we choose $\mathbf{x}^k = \mathbf{x}^*$ then because \mathbf{x}^* satisfies Equation 5, Equation 6 also holds. Conversely, if \mathbf{x}^* is a minimizer of problem (3) at \mathbf{x}^* then \mathbf{x}^* also satisfies Equation 5.

2.1 Intermediate results and gap certificate

We prove several lemmas and we exhibit a gap certificate that provide a bound on the difference of the objective value along the iterations to the optimal objective value.

As one may remark our algorithm is very similar to a conditional gradient algorithm *a.k.a* Frank-Wolfe algorithm. As such, our proof of convergence of the algorithm will follow similar

lines as those used by Bertsekas. Our convergence results is based on the following proposition and definition given in [Ber99].

Proposition 2.2 [Ber99] *Let $\{\mathbf{x}^k\}$ be a sequence generated by the feasible direction method $\mathbf{x}^{k+1} = \mathbf{x}^k + \alpha_k \Delta \mathbf{x}$ with $\Delta \mathbf{x}^k = \mathbf{s}^k - \mathbf{x}^k$. Assume that $\{\Delta \mathbf{x}^k\}$ is gradient related and that α^k is chosen by the limited minimization or the Armijo rule, then every limit point of $\{\mathbf{x}^k\}$ is a stationary point.*

Definition A sequence $\Delta \mathbf{x}^k$ is said to be gradient related to the sequence \mathbf{x}^k if for any subsequence of $\{\mathbf{x}^k\}_{k \in K}$ that converges to a non-stationary point, the corresponding subsequence $\{\Delta \mathbf{x}^k\}_{k \in K}$ is bounded and satisfies

$$\limsup_{k \rightarrow \infty, k \in K} \nabla F(\mathbf{x}^k)^\top \Delta \mathbf{x}^k < 0$$

Basically, this property says that if a subsequence converges to a non-stationary point, then at the limit point the feasible direction defined by $\Delta \mathbf{x}$ is still a descent direction. Before proving that the sequence defined by $\{\Delta \mathbf{x}^k\}$ is gradient related, we prove useful lemmas.

Lemma 2.3 *For any $\mathbf{x}^k \in \mathcal{P}$, each $\Delta \mathbf{x}^k = \mathbf{s}^k - \mathbf{x}^k$ defines a feasible descent direction.*

Proof By definition, \mathbf{s}^k belongs to the convex set \mathcal{P} . Hence, for any $\alpha^k \in [0, 1]$, \mathbf{x}^{k+1} defines a feasible point. Hence $\Delta \mathbf{x}^k$ is a feasible direction.

Now let us show that it also defines a descent direction. By definition of the minimizer \mathbf{s}^k , we have for all $\mathbf{s} \in \mathcal{P}$

$$\langle \nabla f(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle + g(\mathbf{s}^k) - g(\mathbf{x}^k) \leq \langle \nabla f(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle + g(\mathbf{s}) - g(\mathbf{x}^k)$$

because the above inequality also holds for $\mathbf{s} = \mathbf{x}^k$, we have

$$\langle \nabla f(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle + g(\mathbf{s}^k) - g(\mathbf{x}^k) \leq 0 \quad (7)$$

By convexity of $g(\cdot)$, we have

$$g(\mathbf{s}^k) - g(\mathbf{x}^k) \geq \langle \nabla g(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle$$

which, plugged in equation, 7 leads to

$$\langle \nabla f(\mathbf{x}^k) + \nabla g(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle \leq 0$$

and thus $\langle \nabla F(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle \leq 0$, which proves that $\Delta \mathbf{x}^k$ is a descent direction.

The next lemma provides an interesting feature of our algorithm. Indeed, the lemma states that the difference between the optimal objective value and the current objective value can be easily monitored.

Lemma 2.4 *For all $\mathbf{x}^k \in \mathcal{P}$, the following property holds*

$$\min_{\mathbf{s} \in \mathcal{P}} [\langle \nabla f(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle + g(\mathbf{s}) - g(\mathbf{x}^k)] \leq F(\mathbf{x}^*) - F(\mathbf{x}^k) \leq 0$$

where \mathbf{x}^* is a minimizer of F . In addition, if \mathbf{x}^k does not belong to the set of minimizers of $F(\cdot)$, then the second inequality is strict.

Proof By convexity of f , we have

$$f(\mathbf{x}^*) - f(\mathbf{x}^k) \geq \nabla f(\mathbf{x}^k)^\top (\mathbf{x}^* - \mathbf{x}^k)$$

By adding $g(\mathbf{x}^*) - g(\mathbf{x}^k)$ to both side of the inequality, we obtain

$$F(\mathbf{x}^*) - F(\mathbf{x}^k) \geq \nabla f(\mathbf{x}^k)^\top (\mathbf{x}^* - \mathbf{x}^k) + g(\mathbf{x}^*) - g(\mathbf{x}^k)$$

and because \mathbf{x}^* is a minimizer of F , we also have $0 \geq F(\mathbf{x}^*) - F(\mathbf{x}^k)$ Hence, the following holds

$$\langle \nabla f(\mathbf{x}^k), \mathbf{x}^* - \mathbf{x}^k \rangle + g(\mathbf{x}^*) - g(\mathbf{x}^k) \leq F(\mathbf{x}^*) - F(\mathbf{x}^k) \leq 0$$

and we also have

$$\min_{\mathbf{s} \in \mathcal{P}} \langle \nabla f(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle + g(\mathbf{s}) - g(\mathbf{x}^k) \leq F(\mathbf{x}^*) - F(\mathbf{x}^k) \leq 0$$

which concludes the first part of the proof.

Finally if \mathbf{x}^k is not a minimizer of F , then we naturally have $0 > F(\mathbf{x}^*) - F(\mathbf{x}^k)$.

2.2 Proof of convergence

Now that we have all the pieces of the proof, let us show the key ingredient.

Lemma 2.5 *The sequence $\{\Delta \mathbf{x}^k\}$ of our algorithm is gradient related.*

Proof For showing that our direction sequence is gradient related, we have to show that given a subsequence $\{\mathbf{x}^k\}_{k \in K}$ that converges to a non-stationary point $\tilde{\mathbf{x}}$, the sequence $\{\Delta \mathbf{x}^k\}_{k \in K}$ is bounded and that

$$\limsup_{k \rightarrow \infty, k \in K} \nabla F(\mathbf{x}^k)^\top \Delta \mathbf{x}^k < 0$$

Boundedness of the sequence naturally derives from the fact that $\mathbf{s}^k \in \mathcal{P}$, $\mathbf{x}^k \in \mathcal{P}$ and the set \mathcal{P} is bounded.

The second part of the proof starts by showing that

$$\langle \nabla F(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle = \langle \nabla f(\mathbf{x}^k) + \nabla g(\mathbf{x}), \mathbf{s}^k - \mathbf{x}^k \rangle \leq \langle \nabla f(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle + g(\mathbf{s}^k) - g(\mathbf{x}^k)$$

where the last inequality is obtained owing to the convexity of g . Because that inequality holds for the minimizer, it also holds for any vector $\mathbf{s} \in \mathcal{P}$:

$$\langle \nabla F(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle \leq \langle \nabla f(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle + g(\mathbf{s}) - g(\mathbf{x}^k)$$

Taking limit yields to

$$\limsup_{k \rightarrow \infty, k \in K} \langle \nabla F(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle \leq \langle \nabla f(\tilde{\mathbf{x}}), \mathbf{s} - \tilde{\mathbf{x}} \rangle + g(\mathbf{s}) - g(\tilde{\mathbf{x}})$$

for all $\mathbf{s} \in \mathcal{P}$. As such, this inequality also holds for the minimizer

$$\limsup_{k \rightarrow \infty, k \in K} \langle \nabla F(\mathbf{x}^k), \mathbf{s}^k - \mathbf{x}^k \rangle \leq \min_{\mathbf{s} \in \mathcal{P}} \langle \nabla f(\tilde{\mathbf{x}}), \mathbf{s} - \tilde{\mathbf{x}} \rangle + g(\mathbf{s}) - g(\tilde{\mathbf{x}})$$

Now, since $\tilde{\mathbf{x}}$ is not stationary, it is not optimal and it does not belong to the minimizer of F , hence according to the above lemma 2.4,

$$\min_{\mathbf{s} \in \mathcal{P}} \langle \nabla f(\tilde{\mathbf{x}}), \mathbf{s} - \tilde{\mathbf{x}} \rangle + g(\mathbf{s}) - g(\tilde{\mathbf{x}}) < 0$$

which concludes the proof.

This latter lemma proves that our direction sequence is gradient related, thus proposition 2.2 applies.

2.3 Rate of convergence

We can show that the objective value $F(\mathbf{x}_k)$ converges towards $F(\mathbf{x}^*)$ in a linear rate if we have some additional smoothness condition of $F(\cdot)$. We can easily prove this statement by following the steps proposed by Jaggi et al. [Jag13] for the conditional gradient algorithm.

We make the hypothesis that there exists a constant C_F so that for any $\mathbf{x}, \mathbf{y} \in \mathcal{P}$ and any $\alpha \in [0, 1]$, so that the inequality

$$F((1 - \alpha)\mathbf{x} + \alpha\mathbf{y}) \leq F(\mathbf{x}) + \alpha \nabla F(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{C_F}{2} \alpha^2$$

holds.

Based on this inequality, for a sequence $\{\mathbf{x}_k\}$ obtained from the generalized conditional gradient algorithm we have

$$F(\mathbf{x}_{k+1}) - F(\mathbf{x}^*) \leq F(\mathbf{x}_k) - F(\mathbf{x}^*) + \alpha_k \nabla F(\mathbf{x}_k)^\top (\mathbf{s}_k - \mathbf{x}_k) + \frac{C_F}{2} \alpha_k^2$$

Let us denote as $h(\mathbf{x}_k) = F(\mathbf{x}_k) - F(\mathbf{x}^*)$, now by adding to both sides of the inequality $\alpha_k[g(\mathbf{s}_k) - g(\mathbf{x}_k)]$ we have

$$\begin{aligned} & h(\mathbf{x}_{k+1}) + \alpha_k[g(\mathbf{s}_k) - g(\mathbf{x}_k)] \\ & \leq h(\mathbf{x}_k) + \alpha_k[\nabla f(\mathbf{x}_k)^\top(\mathbf{x}^* - \mathbf{x}_k) + g(\mathbf{x}^*) - g(\mathbf{x}_k)] + \alpha_k \nabla g(\mathbf{x}_k)^\top(\mathbf{s}_k - \mathbf{x}_k) + \frac{C_F}{2} \alpha_k^2 \end{aligned}$$

where the second inequality comes from the definition of the search direction \mathbf{s}_k . Now because, $f(\cdot)$ is convex we have $f(\mathbf{x}^*) - f(\mathbf{x}_k) \geq \nabla f(\mathbf{x}_k)^\top(\mathbf{x}^* - \mathbf{x}_k)$. Thus we have

$$h(\mathbf{x}_{k+1}) + \alpha_k[g(\mathbf{s}_k) - g(\mathbf{x}_k)] \leq (1 - \alpha_k)h(\mathbf{x}_k) + \alpha_k \nabla g(\mathbf{x}_k)^\top(\mathbf{s}_k - \mathbf{x}_k) + \frac{C_F}{2} \alpha_k^2$$

Now, again, owing to the convexity of $g(\cdot)$, we have $0 \geq -g(\mathbf{s}_k) + g(\mathbf{x}_k) \nabla g(\mathbf{x}_k)^\top(\mathbf{s}_k - \mathbf{x}_k)$. Using this fact in the last above inequality leads to

$$h(\mathbf{x}_{k+1}) \leq (1 - \alpha_k)h(\mathbf{x}_k) + \frac{C_F}{2} \alpha_k^2 \quad (8)$$

Based on this result, we can now state the following

Theorem 2.6 *For each $k \geq 1$, the iterates \mathbf{x}_k of Algorithm 1 satisfy*

$$F(\mathbf{x}_k) - F(\mathbf{x}^*) \leq \frac{2C_F}{k+2}$$

Proof The proof stands on Equation 8 and on the same induction as the one used by Jaggi et al [Jag13].

Note that this convergence property would also hold if we choose the step size as $\alpha_k = \frac{2}{k+2}$.

2.4 Related works and discussion

Here, we propose to take advantage of the composite structure of the objective function and of an efficient solver of that partially linearized problem. Note that a result similar to Lemma 2.4, denoted as the surrogate duality gap in [Jag13], exists for the conditional gradient:

$$\min_{\mathbf{s} \in \mathcal{P}} \langle \nabla F(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle \leq F(\mathbf{x}^*) - F(\mathbf{x}^k) \leq 0.$$

Using the convexity of $g(\cdot)$ one can see that

$$\langle \nabla F(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle \leq \langle \nabla f(\mathbf{x}^k), \mathbf{s} - \mathbf{x}^k \rangle + g(\mathbf{s}) - g(\mathbf{x}^k).$$

This means that the bound expressed in Lemma 2.4 is at least as good as the one provided by the classical CG. In addition when $g(\cdot)$ is strictly convex, our bound is strictly better which suggests that our approach provides a better control of the convergence along the iterations.

The approach that is the most related to this method is probably the conditional gradient for composite optimization proposed in [HJN13]. In their work, the authors show that when the constraint can be expressed as a bound on a norm, it is can be more efficient to solve an equivalent regularized version, *i.e.* a composite problem. By solving the equivalent problem, they can benefit from efficient computation for nuclear norm regularization and Total Variation in images. The main difference with our approach is that they focus on potentially nondifferentiable norm based regularization for $g(\cdot)$ whereas our algorithm can be used for any convex and differentiable $g(\cdot)$.

Finally our approach is also closely related to projected gradient descent and its spectral variant [BMR00]. As discussed more in detail in section 3.2, when $g(\cdot)$ is an euclidean distance, the solving problem (2) boils down to a projection onto the convex \mathcal{P} . In practice the method is more general since it can be used for any problem as long as problem (2) can be efficiently solved.

3 Application to machine learning problems and numerical experiments

In this section, we showcase two interesting applications of CGS on difficult machine learning problems. First, we discuss the application of CGS on the regularized optimal transport problem, and then we consider a more general application framework of our conditional gradient splitting algorithm to the widely used elastic-net regularization.

3.1 Regularized optimal transport

Regularized optimal transport have recently been developed as an elegant way to model several classical problems in machine learning and image processing. For instance, it is used for color transfer in images, a problem that consists in matching the colors of a source image to that of a target one [FPPA14, RP14]. It has also been successfully applied to the problem of unsupervised domain adaptation where one seeks to adapt a classifier from a source domain (used for learning) to a target domain (used for testing) [CFT14] and recently it has been also considered as an efficient way to compute Wasserstein barycenters in [CD14].

Regularized optimal transport consists in searching an optimal transportation plan to move one source distribution onto a target one, with particular conditions on the plan. In the particular case where only discrete samples $\{\mathbf{x}_i^s\}_{1 \leq i \leq n_s}$ and $\{\mathbf{x}_i^t\}_{1 \leq i \leq n_t}$ are available for respectively the source and target distributions, we note the corresponding distributions as vectors $\mu^s \in \mathbb{R}^{+n_s}$ and $\mu^t \in \mathbb{R}^{+n_t}$. Usually, μ^s and μ^t are seen as histograms since both belong to the probability simplex. In the optimal transport problem, those distributions are embedded in metric spaces, which allows to define a transport cost. In the discrete case, this metric is given as a matrix \mathbf{C} , for which each pairwise term $C_{i,j}$ measures the cost of transporting the i th component of μ^s onto the j th component of μ^t . OT aims at finding a positive matrix γ which can be seen as a joint probability distribution between the source and target, with marginals μ^s and μ^t . It belongs to the set of doubly stochastic matrices or Birkhoff polytope. The optimal transport plan is the one which minimizes the total cost of the transport of μ^s onto μ^t . The regularization applies on the matrix γ , and aims at favoring particular conditioning of this plan. The corresponding optimization problem reads

$$\gamma_0^\lambda = \underset{\gamma \in \mathcal{P}}{\operatorname{argmin}} \quad \langle \gamma, \mathbf{C} \rangle_F + \lambda \Omega(\gamma), \quad (9)$$

$$\text{s.t. } \gamma \geq 0, \quad \gamma \mathbf{1}_{n_t} = \mu^s, \quad \gamma^\top \mathbf{1}_{n_s} = \mu^t$$

where $\Omega(\cdot)$ is a regularization term. In the particular case where one considers an information theoretic measure on γ , namely the negentropy, this term can be written as $\Omega(\gamma) = \Omega_{\text{IT}}(\gamma) = \sum_{i,j} \gamma(i,j) \log \gamma(i,j)$. **(author?)** [Cut13] proposed an extremely efficient algorithm, which uses the scaling matrix approach of Sinkhorn-Knopp [Kni08] to solve this problem.

Other types of regularizations can be considered. In [CFT14], authors use a group sparse regularization term to prevent elements from different classes to be matched in the target domain. In [FPPA14], γ is regularized such that an estimated positions of the source samples, transported in the target domain, are consistently moved with respect to their initial spatial distribution. It has been applied with success to color transfer between images where pixels are seen as 3D samples. The same approach has been also tested for domain adaptation and 3D shape matching in [FCTR14]. It can be seen a Laplacian regularization, where the regularization term reads

$$\Omega_{\text{Lap}}(\gamma) = \lambda_s \operatorname{Tr}(\mathbf{X}_t^\top \gamma^\top \mathbf{L}_s \gamma \mathbf{X}_t) + \lambda_t \operatorname{Tr}(\mathbf{X}_s^\top \gamma \mathbf{L}_t \gamma^\top \mathbf{X}_s).$$

Here, \mathbf{L}_s and \mathbf{L}_t are the Laplacian matrices and \mathbf{X}_s and \mathbf{X}_t are the matrices of source and target samples positions. The two terms both aim at preserving the shape of the two distributions, with respective regularization parameters λ_s and λ_t .

We consider in this paper a problem which combines the two regularizations:

$$\gamma_0^\lambda = \underset{\gamma}{\operatorname{argmin}} \quad \langle \gamma, \mathbf{C} \rangle_F + \lambda_1 \Omega_{\text{IT}}(\gamma) + \lambda_2 \Omega_{\text{Lap}}(\gamma), \quad (10)$$

$$\text{s.t. } \gamma \geq 0, \quad \gamma \mathbf{1}_{n_t} = \mu^s, \quad \gamma^\top \mathbf{1}_{n_s} = \mu^t$$

This problem is hard to solve, for several reasons: the presence of the entropy related term prevents quadratic programming strategies to be applied, and because the Laplacian matrices are generally dense the objective function can be costly to compute. But this problem fits particularly well into the CGS framework. Indeed, we have a constrained optimization problem where the objective function is smooth and we have at our disposal the efficient algorithm of [Cut13] that is able to solve a partially linearized version of the objective function under the same constraints. In this context, we define $f(\gamma) = \langle \gamma, \mathbf{C} \rangle_F + \lambda_2 \Omega_{\text{Lap}}(\gamma)$ and $g(\gamma) = \lambda_1 \Omega_{\text{IT}}(\gamma)$. According to these definitions, the problem of finding \mathbf{s}_k boils down to

$$\mathbf{s}_k = \underset{\gamma}{\operatorname{argmin}} \quad \langle \gamma, \mathbf{C} + \lambda_2 \nabla \Omega_{\text{Lap}}(\gamma) \rangle_F + \lambda_1 \Omega_{\text{IT}}(\gamma), \quad (11)$$

$$\text{s.t. } \mathbf{s} \geq 0, \quad \gamma \mathbf{1}_{n_t} = \mu^s, \quad \gamma^\top \mathbf{1}_{n_s} = \mu^t$$

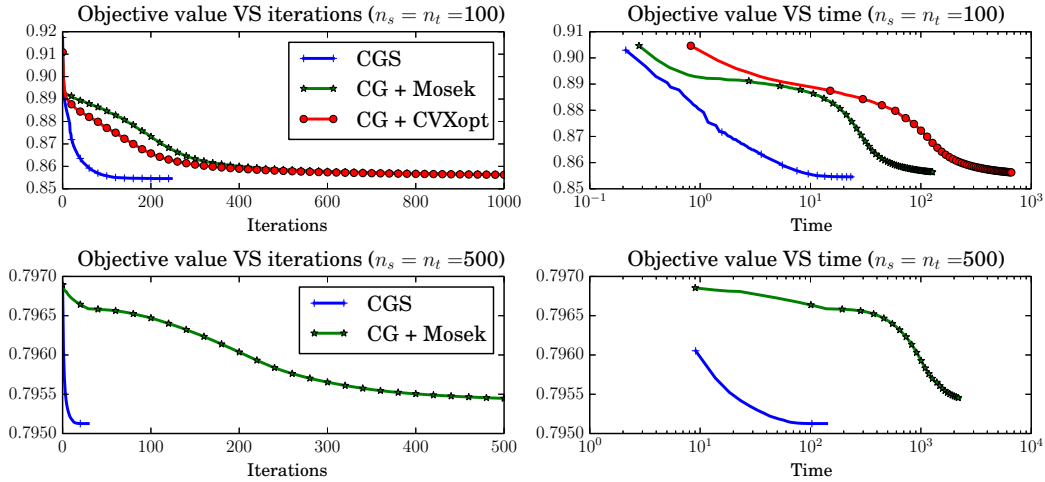


Figure 1: Objective value for a regularized optimal transport problem of size 100x100 (top) and 500x500 (bottom) along the iterations and along time for CGS and CG with different solvers.

for which an efficient solver exists as it is equivalent to the problem addressed in Equation 9 in the particular case of the negentropy regularization [Cut13]. Note that while Ω_{IT} is not differentiable in 0, the Sinkhorn-Knopp algorithm never leads to exact zero coefficients (See Fig; 1 in [Cut13]), hence $g(\cdot)$ is differentiable for all iterations.

We study the performances of our approach on a simulated example similar to the one illustrated in Figure 3.1 of [FPPA14]. Samples from the source and target domains are generated with a noisy cluster structure. In order to keep this structure the symmetric Laplacian regularization is constructed from the graph of 10 nearest neighbors in each domain. The regularization parameters have been selected as to promote a transport of the graph structure with a reasonable regularization ($\lambda_1 = 1.7 \times 10^{-2}$ and $\lambda_2 = 10^3$). The experiments are performed with the dimensions $n_s = n_t = 100$ and $n_s = n_t = 500$, leading to a total number of variables $n_s \times n_t$. In the experiments we compare our method with the conditional gradient algorithm on the exact same problem. In this case, fully linearizing the objective function leads to a linear problem for finding \mathbf{s}_k . Note that while other approaches based on proximal methods have been proposed to solve this kind of optimization problem [PPO14], we do not think we can compare fairly since they are not interior point methods and the intermediate iterations might violate the constraints.

We report in Figure 1 both objective values along iterations (left column) and along the overall computational time (right column) for the two problems. Two different implementations of linear programming solvers were used for comparisons: CVXopt [DV08] and MOSEK [AA00]. In the second examples, the CVXopt experiments were not reported as it took too much time to complete the optimization. As one can see, the CGS outperforms in all cases the CG approaches, both in overall computation time, and also because it reaches better objective values. This difference is also amplified when the size of the problem increases, and can be seen on the last row. Note that the gain in computational time brought by our algorithm is about an order of magnitude better than a CG algorithm using MOSEK.

3.2 Learning with elastic-net regularization

Elastic-net regularization has been introduced by [ZH05] as a way to balance some undesirable behaviour of the ℓ_1 penalty. It has been mainly motivated by the fact that it allows the selection of groups of correlated features and yields consistent prediction and feature selection [DMDVR09]. This regularization is composed of the weighted sum of a ℓ_1 -norm and a squared ℓ_2 -norm regularization. In this work, we want to show that the ℓ_1 norm-constrained version of the elastic-net problem can be efficiently learned by our conditional gradient splitting algorithm. Denote as $\{\mathbf{z}_i, y_i\}_{1 \leq i \leq m}$ a set of observed data, with $\mathbf{z}_i \in \mathbb{R}^n$ being a feature vector and $y_i \in \mathbb{R}$ the target samples. The objective is to learn a linear model $\mathbf{z}^\top \mathbf{x}$ that predicts the target y . For this purpose, we consider the

following constrained elastic-net problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} L(\mathbf{y}, \mathbf{Z}\mathbf{x}) + \lambda \mathbf{x}^\top \mathbf{x} \quad (12)$$

$$\text{s.t. } \|\mathbf{x}\|_1 \leq \tau \quad (13)$$

where $\mathbf{Z} \in \mathbb{R}^{m \times n}$ is the matrix composed by stacking rows of \mathbf{z}_i^\top , \mathbf{y} is the vector of the y_i . and $L(\cdot, \cdot)$ is a differentiable loss function that measures the discrepancy between each coordinate of \mathbf{y} and $\mathbf{Z}\mathbf{x}$. In the context of our conditional gradient splitting algorithm, we will define $f(\mathbf{x})$ as $L(\mathbf{y}, \mathbf{Z}\mathbf{x})$ and $g(\mathbf{x}) = \lambda \mathbf{x}^\top \mathbf{x}$. Accordingly, the step 3 of the algorithm becomes

$$\mathbf{s}^k = \operatorname{argmin}_{\|\mathbf{x}\|_1 \leq \tau} \mathbf{x}^\top \nabla f(\mathbf{x}^k) + \lambda \mathbf{x}^\top \mathbf{x}$$

Interestingly, this problem can be easily solved as it can be shown that the above problem is equivalent to

$$\mathbf{s}^k = \operatorname{argmin}_{\|\mathbf{x}\|_1 \leq \tau} \left\| \mathbf{x} - \left(-\frac{1}{2\lambda} \nabla f(\mathbf{x}^k) \right) \right\|_2^2$$

Hence, the feasible point \mathbf{s}_k is the projection of the scaled negative gradient onto the set defined by the constraint. As such, in this particular case where $g(\mathbf{x})$ is a quadratic term, our algorithm has the flavor of a gradient projection method. The main difference resides in the fact that in our CGS algorithm, it is the negative gradient that is projected onto the constraint set, instead of the point resulting from a step along the negative gradient. Formally, at each iteration, we thus have

$$\mathbf{x}^{k+1} = (1 - \alpha)\mathbf{x}^k + \alpha \Pi_{\|\mathbf{x}\|_1 \leq \tau} \left(-\frac{1}{2\lambda} \nabla f(\mathbf{x}^k) \right) \quad (14)$$

and \mathbf{x}^{k+1} is a linear combination of the current iterate and the projected scaled negative gradient.

This framework can be extended to any constraint set, and we can expect the algorithm to be efficient as long as projection of the set can be computed in a cheap way. In addition the algorithm can be used for any convex and differentiable data fitting term and thus it can be used also for classification with squared hinge loss [Cha07] or logistic regression [KCFH05]. Note however that in these latter cases, the optimal step α can not be computed in a closed form as in a least-square loss context [FPPA14].

In the following, we have illustrated the behaviour of our conditional gradient splitting algorithm compared to classical projected gradient algorithms on toy and real-world classification problems using an elastic-net logistic regression problem. As such, we have considered the limited-memory projected quasi-newton (PQN) method [SBFM09] and the spectral projected gradient (SPG) method [BMR00] both implemented by Schmidt. In our comparisons, we have also included the original conditional gradient algorithm as well as an heuristic conditional gradient splitting with step α^k set as $\alpha^k = \frac{2}{k+2}$. For all algorithms, we have used a monotone armijo rule as a linesearch algorithm and the stopping criterion is based on the fixed point property of a minimizer of (13) (also used by Schmidt)

$$\left\| \Pi_{\|\mathbf{x}\|_1 \leq \tau}(\mathbf{x} - \nabla F(\mathbf{x})) - \mathbf{x} \right\|_\infty \leq \varepsilon$$

In our experiments, we have set $\varepsilon = 10^{-5}$ and we have also set the maximal number of iterations to 10000 for all algorithms. Our objective in these experiments is essentially to show that our conditional gradient splitting is as efficient as other competitors.

The toy problem is the same as the one used by [OTJ10]. The task is a binary classification problem in \mathbb{R}^d . Among these d variables, only T of them define a subspace of \mathbb{R}^d in which classes can be discriminated. For these T relevant variables, the two classes follow a Gaussian pdf with means respectively μ and $-\mu$ and covariance matrices randomly drawn from a Wishart distribution. μ has been randomly drawn from $\{-1, +1\}^T$. The other $d - T$ non-relevant variables follow an *i.i.d* Gaussian probability distribution with zero mean and unit variance for both classes. We have sampled N examples and used 80% of them for training and the rest for the testing. Before learning, the training set has been normalized to zero mean and unit variance and test set has been rescaled accordingly. The hyperparameters λ and τ have been roughly set so as to maximize the performance on the test set. We have chosen to initialize all algorithms with the zero vector.

Figure 2 presents some examples of how the optimality condition of each method evolves with respect to time for different settings of number of examples, variables and number of relevant variables. Independently of the settings, we can note that the conditional gradient algorithm performs very poorly and is not efficient at all compared to all other algorithms. Compared to a projected quasi-newton, our conditional gradient splitting algorithm is far more efficient, and on all the settings it converges faster. Finally, it appears that our algorithm performs on par with the spectral projected gradient algorithm, as it is sometimes faster and in other cases slower. This is a very interesting feature given the simplicity of the algorithm steps. In addition, we can note the nice behaviour of our CGS algorithm with empirical steps which is globally less efficient than CGS with linesearch and the SPG algorithms but provide better improvements of the optimality condition in the first iterations.

For illustrating the algorithm behaviour on real datasets, we have considered two bioinformatic problems for which few examples are available while the number of feature is large: the *colon* and *lymphoma* datasets. We have used the same experimental setting as for the toy dataset. Figure 3 reports typical examples of convergence behaviour. We can note again that our CGS algorithm is slightly more efficient than the spectral projected gradient algorithm and more efficient than the limited-memory projected quasi-newton algorithm. Interestingly, in these real problems, the CGS algorithm with fixed step is the most efficient one for reaching rough optimality conditions (of the order of 10^{-2}).

Acknowledgments

This work was partly funded by the CNRS PEPS Fascido program under the Topase project.

References

- [AA00] Erling Andersen and Knud Andersen. The mosek interior point optimizer for linear programming: An implementation of the homogeneous algorithm. In *High Performance Optimization*, volume 33 of *Applied Optimization*, pages 197–232. Springer US, 2000.
- [Ber99] Dimitri P Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.
- [BLM05] Kristian Bredies, Dirk Lorenz, and Peter Maass. *Equivalence of a generalized conditional gradient method and the method of surrogate functionals*. Zentrum für Technomathematik, 2005.
- [BMR00] Ernesto G Birgin, José Mario Martínez, and Marcos Raydan. Nonmonotone spectral projected gradient methods on convex sets. *SIAM Journal on Optimization*, 10(4):1196–1211, 2000.
- [BNO03] D. Bertsekas, A. Nedic, and A. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, 2003.
- [CD14] Marco Cuturi and Arnaud Doucet. Fast computation of wasserstein barycenters. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, jun 2014.
- [CFT14] Nicolas Courty, Rémi Flamary, and Devis Tuia. Domain adaptation with regularized optimal transport. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, 2014.
- [CFTR15] Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *CoRR*, abs/1507.00504, 2015.
- [Cha07] Olivier Chapelle. Training a support vector machine in the primal. *Neural Computation*, 19(5):1155–1178, 2007.
- [Cut13] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transportation. In *NIPS*, pages 2292–2300. 2013.
- [DMDVR09] Christine De Mol, Ernesto De Vito, and Lorenzo Rosasco. Elastic-net regularization in learning theory. *Journal of Complexity*, 25(2):201–230, 2009.
- [DV08] Joachim Dahl and Lieven Vandenbergh. CVXOPT: A python package for convex optimization, 2008.

- [FCTR14] Rémi Flamary, Nicolas Courty, Devis Tuia, and Alain Rakotomamonjy. Optimal transport with laplacian regularization: Applications to domain adaptation and shape matching. In *NIPS Workshop on Optimal Transport and Machine Learning OTML*, 2014.
- [FPPA14] Sira Ferradans, Nicolas Papadakis, Gabriel Peyré, and Jean-François Aujol. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences*, 7(3), 2014.
- [HJN13] Zaid Harchaoui, Anatoli Juditsky, and Arkadi Nemirovski. Conditional gradient algorithms for norm-regularized smooth convex optimization. *Mathematical Programming*, pages 1–38, 2013.
- [Jag13] Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 427–435, 2013.
- [KCFH05] Balaji Krishnapuram, Lawrence Carin, Mario AT Figueiredo, and Alexander J Hartemink. Sparse multinomial logistic regression: Fast algorithms and generalization bounds. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(6):957–968, 2005.
- [Kni08] Philip Knight. The sinkhorn-knopp algorithm: Convergence and applications. *SIAM J. Matrix Anal. Appl.*, 30(1):261–275, March 2008.
- [OTJ10] G. Obozinski, B. Taskar, and M.I. Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. *Statistics and Computing*, 20:231–252, 2010.
- [PPO14] N. Papadakis, G. Peyré, and E. Oudet. Optimal Transport with Proximal Splitting. *SIAM Journal on Imaging Sciences*, 7(1):212–238, 2014.
- [RP14] Julien Rabin and Nicolas Papadakis. Non-convex relaxation of optimal transport for color transfer. In *NIPS Workshop on Optimal Transport for Machine Learning*, 2014.
- [SBFM09] Mark Schmidt, Ewout Berg, Michael Friedlander, and Kevin Murphy. Optimizing costly functions with simple constraints: A limited-memory projected quasi-newton algorithm. In *AISTATS'09*, page None, 2009.
- [ZH05] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.




Figure 2: Examples of evolution of the optimality condition for three different learning setting. (left) highly-sparse and very few examples. (middle) sparse and few examples. (right) sparse with reasonable ratio of examples over variables.

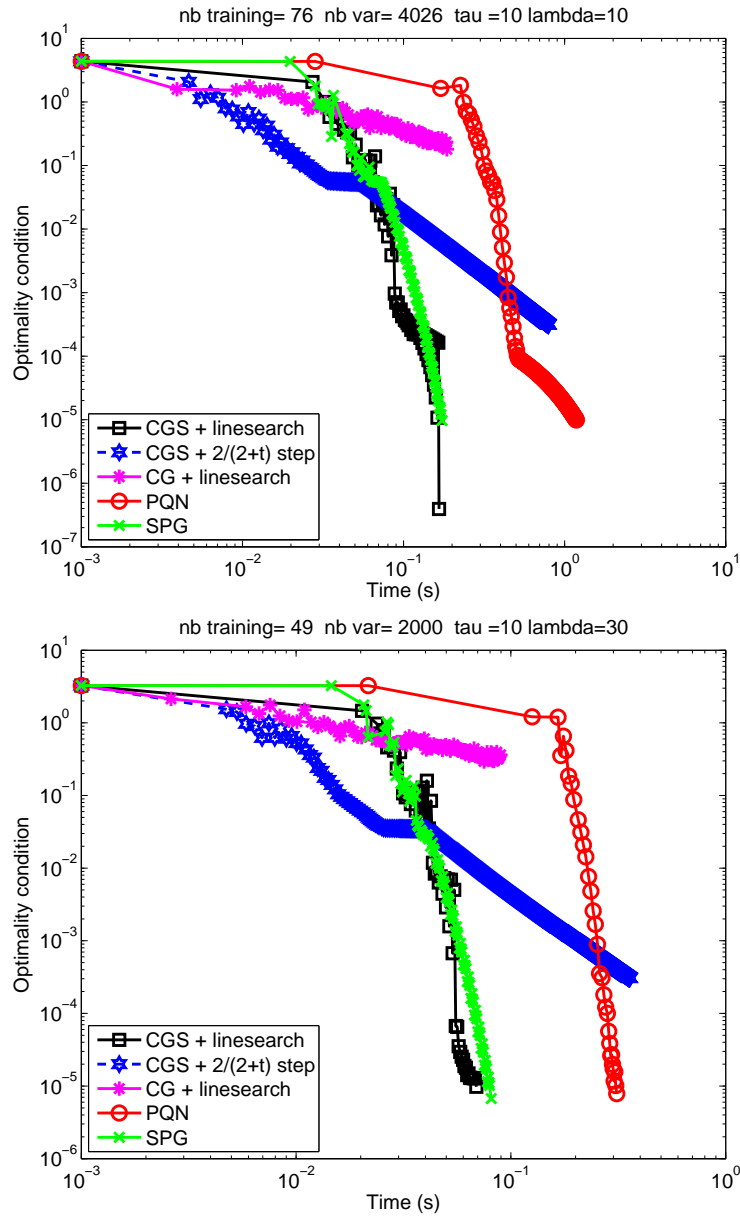


Figure 3: Optimality conditions evolving curves on the (left) *colon* and (right) *lymphoma* datasets.