

Chapter 23 Algorithms for Constrained Optimization

An Introduction to Optimization
Spring, 2015

Wei-Ta Chu

Projections

- ▶ The optimization algorithms described in Part II have the general form $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha_k \mathbf{d}^{(k)}$, where $\mathbf{d}^{(k)}$ is typically a function of $\nabla f(\mathbf{x}^{(k)})$. The value of $\mathbf{x}^{(k)}$ is not constrained to lie inside any particular set.

- ▶ Consider the optimization problem

$$\begin{aligned} & \text{minimize } f(\mathbf{x}) \\ & \text{subject to } \mathbf{x} \in \Omega \end{aligned}$$

- ▶ Modify the algorithm we already know to take constraints into account.

Projections

- ▶ If $\mathbf{x}^{(k)} + \alpha_k \mathbf{d}^{(k)}$ is in Ω , then we set form $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha_k \mathbf{d}^{(k)}$ as usual.
- ▶ If, on the other hand, $\mathbf{x}^{(k)} + \alpha_k \mathbf{d}^{(k)}$ is not in Ω , then we “project” it back into Ω before setting $\mathbf{x}^{(k+1)}$.
- ▶ Consider the case where the constraint set $\Omega \in R^n$ is given by

$$\Omega = \{\mathbf{x}: l_i \leq x_i \leq u_i, \quad i = 1, \dots, n\}$$

- ▶ In this case, Ω is a “box” in R^n ; for this reason, this form of Ω is called a *box constraint*.

Projections

- ▶ Given a point $\mathbf{x} \in R^n$, define $\mathbf{y} = \Pi[\mathbf{x}] \in R^n$ by

$$y_i = \min\{u_i, \max\{l_i, x_i\}\} = \begin{cases} u_i & \text{if } x_i > u_i \\ x_i & \text{if } l_i \leq x_i \leq u_i \\ l_i & \text{if } x_i < l_i \end{cases}$$

- ▶ The point $\Pi[\mathbf{x}]$ is called the projection of \mathbf{x} onto Ω . Note that $\Pi[\mathbf{x}]$ is actually the “closest” point in Ω to \mathbf{x} . Using the projection operator Π , we can modify the previous unconstrained algorithm as follows:

$$\mathbf{x}^{(k+1)} = \Pi[\mathbf{x}^{(k)} + \alpha_k \mathbf{d}^{(k)}]$$

- ▶ Note that the iterates $\mathbf{x}^{(k)}$ now all lie inside Ω . We call the algorithm above a *projection algorithm*.

Projections

- ▶ In the more general case, we can define the projection onto Ω :

$$\Pi[\mathbf{x}] = \arg \min_{\mathbf{z} \in \Omega} \|\mathbf{z} - \mathbf{x}\|$$

- ▶ In this case, $\Pi[\mathbf{x}]$ is again the “closest” point in Ω to \mathbf{x} . This projection operator is well-defined only for certain types of constraint sets: for example, closed convex sets.
- ▶ For some sets Ω , the “arg min” above is not well-defined. If the projection Π is well-defined, we can simply apply the projection algorithm

$$\mathbf{x}^{(k+1)} = \Pi[\mathbf{x}^{(k)} + \alpha_k \mathbf{d}^{(k)}]$$

Projections

- ▶ Consider the projection method applied specifically to the gradient algorithm. Recall that the vector $-\nabla f(\mathbf{x})$ points in the direction of maximum rate of decrease of f at \mathbf{x} . This was the basis for gradient methods for unconstrained optimization, which have the form $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \alpha_k \nabla f(\mathbf{x}^{(k)})$, where α_k is the step size. The choice of the step size depends on the particular gradient algorithm. For example, recall that in the steepest descent algorithm, $\alpha_k = \arg \min_{\alpha \geq 0} f(\mathbf{x}^{(k)} - \alpha \nabla f(\mathbf{x}^{(k)}))$.

Projections

- ▶ The projected version of the gradient algorithm has the form

$$\mathbf{x}^{(k+1)} = \Pi[\mathbf{x}^{(k)} - \alpha_k \nabla f(\mathbf{x}^{(k)})]$$

- ▶ We refer to the above as the *projected gradient algorithm*.

Example

- ▶ Consider the problem

$$\begin{aligned} & \text{minimize } \frac{1}{2} \mathbf{x}^T \mathbf{Q} \mathbf{x} \\ & \text{subject to } \|\mathbf{x}\|^2 = 1 \end{aligned}$$

where $\mathbf{Q} = \mathbf{Q}^T > 0$. Suppose that we apply a *fixed-step-size projected gradient algorithm* to this problem.

- ▶ a. Derive a formula for the update equation for the algorithm. You may assume that the argument in the projection operator to obtain $\mathbf{x}^{(k)}$ is never zero.

Example

- ▶ The projection operator in this case simply maps any vector to the closest point on the unit circle. Therefore, the projection operator is given by $\Pi[\mathbf{x}] = \mathbf{x}/\|\mathbf{x}\|$, provided that $\mathbf{x} \neq \mathbf{0}$. The update equation is

$$\mathbf{x}^{(k+1)} = \beta_k (\mathbf{x}^{(k)} - \alpha \mathbf{Q} \mathbf{x}^{(k)}) = \beta_k (\mathbf{I} - \alpha \mathbf{Q}) \mathbf{x}^{(k)}$$

where $\beta_k = 1/\|(\mathbf{I} - \alpha \mathbf{Q}) \mathbf{x}^{(k)}\|$ (i.e., it is whatever constant scaling is needed to make $\mathbf{x}^{(k+1)}$ have unit norm).

Example

- ▶ b. It is possible for the algorithm not to converge to an optimal solution even if the step size $\alpha > 0$ is taken to be arbitrarily small?
- ▶ If we start with $\mathbf{x}^{(0)}$ being an eigenvector of \mathbf{Q} , then $\mathbf{x}^{(k)} = \mathbf{x}^{(0)}$ for all k . Therefore, if the corresponding eigenvalue is not the smallest, then clearly the algorithm is stuck at a point that is not optimal.

Example

- ▶ c. Show that for $0 < \alpha < 1/\lambda_{max}$ (where λ_{max} is the largest eigenvalue of \mathbf{Q}), the fixed-step-size projected gradient algorithm (with step size α) converges to an optimal solution, provided that $\mathbf{x}^{(0)}$ is not orthogonal to the eigenvectors of \mathbf{Q} corresponding to the smallest eigenvalue. (Assume that the eigenvalues are distinct.)

Example

- ▶ We have

$$\begin{aligned}\mathbf{x}^{(k+1)} &= \beta_k (\mathbf{I} - \alpha \mathbf{Q}) \mathbf{x}^{(k)} \\ &= \beta_k (\mathbf{I} - \alpha \mathbf{Q}) \left(y_1^{(k)} \mathbf{v}_1 + \cdots + y_n^{(k)} \mathbf{v}_n \right) \\ &= \beta_k \left(y_1^{(k)} (\mathbf{I} - \alpha \mathbf{Q}) \mathbf{v}_1 + \cdots + y_n^{(k)} (\mathbf{I} - \alpha \mathbf{Q}) \mathbf{v}_n \right)\end{aligned}$$

- ▶ But $(\mathbf{I} - \alpha \mathbf{Q}) \mathbf{v}_i = (1 - \alpha \lambda_i) \mathbf{v}_i$, where λ_i is the eigenvalue corresponding to \mathbf{v}_i . Hence,

$$\mathbf{x}^{k+1} = \beta_k \left(y_1^{(k)} (1 - \alpha \lambda_1) \mathbf{v}_1 + \cdots + y_n^{(k)} (1 - \alpha \lambda_n) \mathbf{v}_n \right)$$

which means that $y_i^{(k+1)} = \beta_k y_i^{(k)} (1 - \alpha \lambda_i)$.

Example

- ▶ In other words, $y_i^{(k)} = \beta^{(k)} y_i^{(0)} (1 - \alpha \lambda_i)^k$, where $\beta^{(k)} = \prod_{i=0}^{k-1} \beta_k$. We rewrite $\mathbf{x}^{(k)}$ as

$$\mathbf{x}^{(k)} = \sum_{i=1}^n y_i^{(k)} \mathbf{v}_i = y_1^{(k)} \left(\mathbf{v}_1 + \sum_{i=2}^n \frac{y_i^{(k)}}{y_1^{(k)}} \mathbf{v}_i \right)$$

- ▶ Assuming that $y_1^{(0)} \neq 0$, we obtain

$$\frac{y_i^{(k)}}{y_1^{(k)}} = \frac{y_i^{(0)} (1 - \alpha \lambda_i)^k}{y_1^{(0)} (1 - \alpha \lambda_1)^k} = \frac{y_i^{(0)}}{y_1^{(0)}} \left(\frac{1 - \alpha \lambda_i}{1 - \alpha \lambda_1} \right)^k$$

Example

- ▶ Using the fact that $\frac{1-\alpha\lambda_i}{1-\alpha\lambda_1} < 1$ (because the $\lambda_i > \lambda_1$ for $i > 1$ and $\alpha < 1/\lambda_{max}$), we deduce that

$$\frac{y_i^{(k)}}{y_1^{(k)}} \rightarrow 0,$$

which implies that $\mathbf{x}^{(k)} \rightarrow \mathbf{v}_1$, as required.

Projected Gradient Methods with Linear Constraints

- ▶ Consider optimization problems of the form

$$\begin{aligned} & \text{minimize } f(\mathbf{x}) \\ & \text{subject to } \mathbf{Ax} = \mathbf{b} \end{aligned}$$

where $f: R^n \rightarrow R$, $\mathbf{A} \in R^{m \times n}$, $m < n$, $\text{rank } \mathbf{A} = m$, $\mathbf{b} \in R^m$

- ▶ We assume throughout that $f \in C^1$.
- ▶ The specific structure of the constraint set allows us to compute the projection operator Π using the *orthogonal projector*. Specifically, $\Pi[\mathbf{x}]$ can be defined using the orthogonal projector matrix \mathbf{P} given by

$$\mathbf{P} = \mathbf{I}_n - \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}$$

Projected Gradient Methods with Linear Constraints

- ▶ Two important properties of the orthogonal projector \mathbf{P} that we use in this section are
 - ▶ $\mathbf{P} = \mathbf{P}^T$
 - ▶ $\mathbf{P}^2 = \mathbf{P}$
- ▶ Lemma 23.1. Let $\mathbf{v} \in R^n$. Then, $\mathbf{P}\mathbf{v} = \mathbf{0}$ if and only if $\mathbf{v} \in \mathcal{R}(\mathbf{A}^T)$. In other words, $\mathcal{N}(\mathbf{P}) = \mathcal{R}(\mathbf{A}^T)$. Moreover, $\mathbf{A}\mathbf{v} = \mathbf{0}$ if and only if $\mathbf{v} \in \mathcal{R}(\mathbf{P})$; that is, $\mathcal{N}(\mathbf{A}) = \mathcal{R}(\mathbf{P})$.

Projected Gradient Methods with Linear Constraints

- ▶ Recall that in unconstrained optimization, the first-order necessary condition for a point \mathbf{x}^* to be a local minimizer is $\nabla f(\mathbf{x}^*) = 0$.
- ▶ In optimization problems with equality constraints, the Lagrange condition plays the role of the first-order necessary condition.
- ▶ When the constraint set takes the form $\{\mathbf{x}: \mathbf{A}\mathbf{x} = \mathbf{b}\}$, the Lagrange condition can be written as $\mathbf{P}\nabla f(\mathbf{x}^*) = 0$.

Projected Gradient Methods with Linear Constraints

- ▶ Proposition 23.1. Let $\mathbf{x}^* \in R^n$ be a feasible point. Then $\mathbf{P}\nabla f(\mathbf{x}^*) = \mathbf{0}$ if and only if \mathbf{x}^* satisfies the Lagrange condition.
- ▶ Proof. By Lemma 23.1, $\mathbf{P}\nabla f(\mathbf{x}^*) = \mathbf{0}$ if and only if we have $\nabla f(\mathbf{x}^*) \in \mathcal{R}(\mathbf{A}^T)$. This is equivalent to the condition that there exists $\boldsymbol{\lambda}^* \in R^m$ such that $\nabla f(\mathbf{x}^*) + \mathbf{A}^T \boldsymbol{\lambda}^* = \mathbf{0}$, which together with the feasibility equation $\mathbf{A}\mathbf{x} = \mathbf{b}$, constitutes the Lagrange condition.

Projected Gradient Methods with Linear Constraints

- ▶ Recall that the projected gradient algorithm has the form

$$\mathbf{x}^{(k+1)} = \Pi[\mathbf{x}^{(k)} - \alpha_k \nabla f(\mathbf{x}^{(k)})]$$

- ▶ For the case where the constraints are linear, it turns out that we can express the projection Π in terms of the matrix \mathbf{P} as follows:

$$\Pi[\mathbf{x}^{(k)} - \alpha_k \nabla f(\mathbf{x}^{(k)})] = \mathbf{x}^{(k)} - \alpha_k \mathbf{P} \nabla f(\mathbf{x}^{(k)})$$

assuming that $\mathbf{x}^k \in \Omega$.

Projected Gradient Methods with Linear Constraints

- ▶ In our constrained optimization problem, the vector $-\nabla f(\mathbf{x})$ is not necessarily a feasible direction. In other words, if $\mathbf{x}^{(k)}$ is a feasible point and we apply the algorithm $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \alpha_k \nabla f(\mathbf{x}^{(k)})$, then $\mathbf{x}^{(k+1)}$ need not be feasible. This problem can be overcome by replacing $-\nabla f(\mathbf{x}^{(k)})$ by a vector that points in a feasible direction.
- ▶ Note that the set of feasible directions is simply the nullspace $\mathcal{N}(\mathbf{A})$ of the matrix \mathbf{A} . Therefore, we should first project the vector $-\nabla f(\mathbf{x})$ onto $\mathcal{N}(\mathbf{A})$. This projection is equivalent to multiplication by the matrix \mathbf{P} .

Projected Gradient Methods with Linear Constraints

- ▶ In summary, in the projection gradient algorithm, we update $\mathbf{x}^{(k)}$ according to the equation

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \alpha_k \mathbf{P} \nabla f(\mathbf{x}^{(k)})$$

Projected Gradient Methods with Linear Constraints

- ▶ Proposition 23.2. In a projected gradient algorithm, if $\mathbf{x}^{(0)}$ is feasible, then each $\mathbf{x}^{(k)}$ is feasible; that is, for each $k \geq 0$, $\mathbf{Ax}^{(k)} = \mathbf{b}$.
- ▶ Proof. We proceed by induction. The result holds for $k = 0$ by assumption. Suppose now that $\mathbf{Ax}^{(k)} = \mathbf{b}$. We now show that $\mathbf{Ax}^{(k+1)} = \mathbf{b}$. Observe that $\mathbf{P}\nabla f(\mathbf{x}^{(k)}) \in \mathcal{N}(\mathbf{A})$. Therefore,

$$\begin{aligned}\mathbf{Ax}^{(k+1)} &= \mathbf{A}(\mathbf{x}^{(k)} - \alpha_k \mathbf{P}\nabla f(\mathbf{x}^{(k)})) \\ &= \mathbf{Ax}^{(k)} - \alpha_k \mathbf{AP}\nabla f(\mathbf{x}^{(k)}) \\ &= \mathbf{b}\end{aligned}$$

Projected Gradient Methods with Linear Constraints

- ▶ The projected gradient algorithm updates $\mathbf{x}^{(k)}$ in the direction of $-\mathbf{P}\nabla f(\mathbf{x}^{(k)})$. This vector points in the direction of maximum rate of decrease of f at $\mathbf{x}^{(k)}$ along the surface defined by $\mathbf{A}\mathbf{x} = \mathbf{b}$, as described in the following argument.
- ▶ Let \mathbf{x} be any feasible point and \mathbf{d} a feasible direction such that $\|\mathbf{d}\| = 1$. The rate of increase of f at \mathbf{x} in the direction \mathbf{d} is $\langle \nabla f(\mathbf{x}), \mathbf{d} \rangle$. Next, we note that because \mathbf{d} is a feasible direction, it lies in $\mathcal{N}(\mathbf{A})$ and hence by Lemma 23.1, we have $\mathbf{d} \in \mathcal{R}(\mathbf{P}) = \mathcal{R}(\mathbf{P}^T)$. So, there exists \mathbf{v} such that $\mathbf{d} = \mathbf{P}\mathbf{v}$.

Projected Gradient Methods with Linear Constraints

▶ Hence,

$$\langle \nabla f(\mathbf{x}), \mathbf{d} \rangle = \langle \nabla f(\mathbf{x}), \mathbf{P}^T \mathbf{v} \rangle = \langle \mathbf{P} \nabla f(\mathbf{x}), \mathbf{v} \rangle$$

▶ By the Cauchy-Schwarz inequality,

$$\langle \mathbf{P} \nabla f(\mathbf{x}), \mathbf{v} \rangle \leq \|\mathbf{P} \nabla f(\mathbf{x})\| \|\mathbf{v}\|$$

with equality if and only if the direction of \mathbf{v} is parallel with the direction of $\mathbf{P} \nabla f(\mathbf{x})$. Therefore, the vector $-\mathbf{P} \nabla f(\mathbf{x})$ points in the direction of maximum rate of decrease of f at \mathbf{x} among all feasible directions.

Projected Gradient Methods with Linear Constraints

- ▶ Suppose that we have a starting point $\mathbf{x}^{(0)}$, which we assume is feasible; that is, $\mathbf{A}\mathbf{x}^{(0)} = \mathbf{b}$. Consider the point $\mathbf{x} = \mathbf{x}^{(0)} - \alpha\mathbf{P}\nabla f(\mathbf{x}^{(0)})$, where $\alpha \in R$. As usual, the scalar α is called the step size.
- ▶ By the discussion above, \mathbf{x} is also a feasible point. Using a Taylor series expansion of f about $\mathbf{x}^{(0)}$ and the fact that $\mathbf{P} = \mathbf{P}^2 = \mathbf{P}^T\mathbf{P}$, we get

$$\begin{aligned} & f(\mathbf{x}^{(0)} - \alpha\mathbf{P}\nabla f(\mathbf{x}^{(0)})) \\ &= f(\mathbf{x}^{(0)}) - \alpha\nabla f(\mathbf{x}^{(0)})^T \mathbf{P}\nabla f(\mathbf{x}^{(0)}) + o(\alpha) \\ &= f(\mathbf{x}^{(0)}) - \alpha\|\mathbf{P}\nabla f(\mathbf{x}^{(0)})\|^2 + o(\alpha) \end{aligned}$$

Projected Gradient Methods with Linear Constraints

- ▶ Thus, if $\mathbf{P}\nabla f(\mathbf{x}^{(0)}) \neq 0$, that is, $\mathbf{x}^{(0)}$ does not satisfy the Lagrange condition, then we can choose an α sufficiently small such that $f(\mathbf{x}) < f(\mathbf{x}^{(0)})$, which means that $\mathbf{x} = \mathbf{x}^{(0)} - \alpha\mathbf{P}\nabla f(\mathbf{x}^{(0)})$ is an improvement over $\mathbf{x}^{(0)}$.

This is the basis for the projected gradient algorithm $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \alpha_k\mathbf{P}\nabla f(\mathbf{x}^{(k)})$, where the initial point $\mathbf{x}^{(0)}$ satisfies $\mathbf{A}\mathbf{x}^{(0)} = \mathbf{b}$ and α_k is some step size.

Projected Gradient Methods with Linear Constraints

- ▶ A well-known variant: *projected steepest descent algorithm*, where α_k is given by

$$\alpha_k = \arg \min_{\alpha \geq 0} f(\mathbf{x}^{(k)} - \alpha \mathbf{P} \nabla f(\mathbf{x}^{(k)}))$$

- ▶ Theorem 23.1. If $\{\mathbf{x}^{(k)}\}$ is the sequence of points generated by the projected steepest descent algorithm and if $\mathbf{P} \nabla f(\mathbf{x}^{(0)}) \neq 0$, then $f(\mathbf{x}^{(k+1)}) < f(\mathbf{x}^{(k)})$.

Projected Gradient Methods with Linear Constraints

- ▶ If for some k , we have $\mathbf{P}\nabla f(\mathbf{x}^{(k)}) = \mathbf{0}$, then by Proposition 23.1 the point $\mathbf{x}^{(k)}$ satisfies the Lagrange condition. This condition can be used as a stopping criterion for the algorithm. Note that in this case, $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)}$.
- ▶ For the case where f is a convex function, the condition $\mathbf{P}\nabla f(\mathbf{x}^{(k)}) = \mathbf{0}$ is, in fact, equivalent to $\mathbf{x}^{(k)}$ being a global minimizer of f over the constraint set $\{\mathbf{x}: \mathbf{A}\mathbf{x} = \mathbf{b}\}$.

Projected Gradient Methods with Linear Constraints

- ▶ Proposition 23.3. The point $\mathbf{x}^* \in R^n$ is a global minimizer of a convex function f over $\{\mathbf{x}: \mathbf{A}\mathbf{x} = \mathbf{b}\}$ if and only if $\mathbf{P}\nabla f(\mathbf{x}^*) = \mathbf{0}$.
- ▶ Proof. We first write $\mathbf{h}(\mathbf{x}) = \mathbf{A}\mathbf{x} - \mathbf{b}$. Then, the constraints can be written as $\mathbf{h}(\mathbf{x}) = \mathbf{0}$, and the problem is of the form considered in earlier chapters.
- ▶ Note that $D\mathbf{h}(\mathbf{x}) = \mathbf{A}$. Hence, $\mathbf{x}^* \in R^n$ is a global minimizer of f if and only if the Lagrange condition holds (see Theorem 22.8). By Proposition 23.1, this is true if and only if $\mathbf{P}\nabla f(\mathbf{x}^*) = \mathbf{0}$, and this completes the proof.

Lagrangian Algorithm for Equality Constraints

- ▶ The basic idea is to use gradient algorithms to update simultaneously the decision variable and Lagrange multiplier vector.

- ▶ Consider the following problem

$$\begin{aligned} & \text{minimize } f(\mathbf{x}) \\ & \text{subject to } \mathbf{h}(\mathbf{x}) = \mathbf{0} \end{aligned}$$

- ▶ where $\mathbf{h}: R^n \rightarrow R^m$. Recall that for this problem the Lagrangian function is given by

$$l(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) + \boldsymbol{\lambda}^T \mathbf{h}(\mathbf{x})$$

- ▶ Assume that $f, \mathbf{h} \in C^2$; as usual, denote the Hessian of the Lagrangian by $L(\mathbf{x}, \boldsymbol{\lambda})$

Lagrangian Algorithm for Equality Constraints

- ▶ The Lagrangian algorithm for this problem is given by

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \alpha_k \left(\nabla f(\mathbf{x}^{(k)}) + D\mathbf{h}(\mathbf{x}^{(k)})^T \boldsymbol{\lambda}^{(k)} \right)$$

$$\boldsymbol{\lambda}^{(k+1)} = \boldsymbol{\lambda}^{(k)} + \beta_k \mathbf{h}(\mathbf{x}^{(k)})$$

- ▶ Notice that the update equation for $\mathbf{x}^{(k)}$ is a gradient algorithm for minimizing the Lagrangian with respect to its \mathbf{x} argument, and the update equation for $\boldsymbol{\lambda}^{(k)}$ is a gradient algorithm for maximizing the Lagrangian with respect to its $\boldsymbol{\lambda}$ argument.
- ▶ Because only the gradient is used, the method is also called the *first-order Lagrangian algorithm*.

Lagrangian Algorithm for Equality Constraints

- ▶ Lemma 23.2. For the Lagrangian algorithm for updating $\mathbf{x}^{(k)}$ and $\boldsymbol{\lambda}^{(k)}$, the pair $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ is a fixed point if and only if it satisfies the Lagrange condition.
- ▶ Below, we use $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ to denote a pair satisfying the Lagrange condition. Assume that $L(\mathbf{x}^*, \boldsymbol{\lambda}^*) > 0$. Also assume that \mathbf{x}^* is a *regular* point. With these assumptions, we are now ready to state and prove that the algorithm is locally convergent. For simplicity, we will take α_k and β_k to be fixed constants (not depending on k), denoted α and β , respectively.

Lagrangian Algorithm for Equality Constraints

- ▶ Theorem 23.2. For the Lagrangian algorithm for updating $\mathbf{x}^{(k)}$ and $\boldsymbol{\lambda}^{(k)}$, provided that α and β are sufficiently small, there is a neighborhood of $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ such that if the pair $(\mathbf{x}^{(0)}, \boldsymbol{\lambda}^{(0)})$ is in this neighborhood, then the algorithm converges to $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ with at least a linear order of convergence.

Lagrangian Algorithm for Inequality Constraints

- ▶ Consider the following optimization problem with inequality constraints:

$$\begin{aligned} & \text{minimize } f(\mathbf{x}) \\ & \text{subject to } \mathbf{g}(\mathbf{x}) \leq \mathbf{0} \end{aligned}$$

where $\mathbf{g}: R^n \rightarrow R^p$.

- ▶ Recall that for this problem the Lagrangian function is given by

$$l(\mathbf{x}, \boldsymbol{\mu}) = f(\mathbf{x}) + \boldsymbol{\mu}^T \mathbf{g}(\mathbf{x})$$

- ▶ As before, assume that $f, \mathbf{g} \in C^2$; as usual, denote the Hessian of the Lagrangian by $L(\mathbf{x}, \boldsymbol{\mu})$

Lagrangian Algorithm for Inequality Constraints

- ▶ The Lagrangian algorithm for this problem is given by

$$\begin{aligned}\mathbf{x}^{(k+1)} &= \mathbf{x}^{(k)} - \alpha_k \left(\nabla f(\mathbf{x}^{(k)}) + D\mathbf{g}(\mathbf{x}^{(k)})^T \boldsymbol{\mu}^{(k)} \right) \\ \boldsymbol{\mu}^{(k+1)} &= [\boldsymbol{\mu}^{(k)} + \beta_k \mathbf{g}(\mathbf{x}^{(k)})]_+\end{aligned}$$

where $[\cdot]_+ = \max\{\cdot, 0\}$ (applied componentwise).

- ▶ Notice that, as before, the update equation for $\mathbf{x}^{(k)}$ is a gradient algorithm for minimizing the Lagrangian with respect to its \mathbf{x} argument. The update equation for $\boldsymbol{\mu}^{(k)}$ is a *projected gradient algorithm* for maximizing the Lagrangian with respect to its $\boldsymbol{\mu}$ argument. The reason for the projection is that the KKT multiplier vector is required to be nonnegative to satisfy the KKT condition.

Lagrangian Algorithm for Inequality Constraints

- ▶ Lemma 23.3. For the Lagrangian algorithm for updating $\mathbf{x}^{(k)}$ and $\boldsymbol{\mu}^{(k)}$, the pair $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ is a fixed point if and only if it satisfies the KKT condition.
- ▶ As before, we use the notation $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ to denote a pair satisfying the KKT condition. Assume that $L(\mathbf{x}^*, \boldsymbol{\mu}^*) > 0$. Also assume that \mathbf{x}^* is a regular point. With these assumptions, we are now ready to state and prove that the algorithm is locally convergent.

Lagrangian Algorithm for Inequality Constraints

- ▶ As before, we will take α_k and β_k to be fixed constants (not depending on k), denoted α and β , respectively. Our analysis examines the behavior of the algorithm in two phases. In the first phase, the “nonactive” multipliers decrease to zero in finite time and remain at zero thereafter. In the second phase, the $\mathbf{x}^{(k)}$ iterates and the “active” multipliers converge jointly to their respective solutions, with at least a linear order of convergence.

Lagrangian Algorithm for Inequality Constraints

- ▶ Theorem 23.3. For the Lagrangian algorithm for updating $\mathbf{x}^{(k)}$ and $\boldsymbol{\mu}^{(k)}$, provided that α and β are sufficiently small, there is a neighborhood of $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ such that if the pair $(\mathbf{x}^{(0)}, \boldsymbol{\mu}^{(0)})$ is in this neighborhood, then (1) the nonactive multipliers reduce to zero in finite time and remain at zero thereafter and (2) the algorithm converges to $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ with at least a linear order of convergence.

Penalty Methods
