CS 450: Numerical Analysis – Lecture 18

Optimality Conditions and Golden Section Search

Edgar Solomonik

October 29, 2026

Chapter 6: Optimization (Lecture 18, continued).
Nonlinear systems and Broyden's method are in the Chapter 5 notes.

Outline: optimization part of Lecture 18

- Critical points, Hessians, and local optimality conditions
- Newton applied to the gradient and sensitivity of a minimizer
- Golden section search: interval reduction, reuse, and stopping

Optimization (continuous)

Given $f : \mathbb{R}^n \to \mathbb{R}$, solve

$$\min_x \quad \underbrace{f(x)}_{\text{objective function}} \qquad \text{subject to} \qquad \underbrace{g(x) = 0}_{\text{equality constraints}} \quad , \qquad \underbrace{h(x) \leq 0}_{\text{inequality constraints}} \quad .$$

We assume f, g, h differentiable. Today: focus on the unconstrained case (no constraints).

- If f, g, h are affine, the problem is a *linear program.*
- If f is quadratic and g, h are affine, we have a *quadratic program.* Here

$$f(x) = f(0) + \nabla f(0)^T x + \tfrac{1}{2}\, x^T H_f(0)\, x,$$

  where H_f is the *Hessian matrix*,

$$H_f(x) = J_{\nabla f}(x) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1 \partial x_1}(x) & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n}(x) \\ \vdots & & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1}(x) & \cdots & \frac{\partial^2 f}{\partial x_n \partial x_n}(x) \end{bmatrix};$$

  for twice continuously differentiable f, $H_f(x)^T = H_f(x)$.

Unconstrained minima. Assume f is twice continuously differentiable near x^*. If x^* is an unconstrained local minimizer, then $\nabla f(x^*) = 0$ and $H_f(x^*)$ is positive semidefinite (necessary conditions). A point with $\nabla f(x^*) = 0$ is called a *critical point* (or *stationary point*). If the Hessian at a critical point is positive definite, then the point is a strict local minimizer. Thus a positive definite Hessian is sufficient at a critical point, but not necessary. A singular positive semidefinite Hessian is inconclusive: x^* may or may not be a local minimizer. These conditions follow from the second-order Taylor expansion,

$$f(x^* + p) = f(x^*) + \nabla f(x^*)^T p + \tfrac{1}{2} p^T H_f(x^*) p + \text{higher-order terms.} \tag{18.4}$$

At a critical point, a negative eigenvalue of $H_f(x^*)$ gives nearby points with smaller function values by choosing a small p along that eigenvector, ruling out a local minimum. These are local conditions, not tests for global optimality.

Newton's method for optimization. Apply Newton's method to the nonlinear system $\nabla f(x) = 0$:

$$H_f(x_k)s_k = -\nabla f(x_k), \qquad x_{k+1} = x_k + s_k.$$

This finds a critical point, not necessarily a local minimum. For example, for $f(x_1, x_2) = x_1^2 - x_2^2$, Newton applied to the gradient reaches $(0, 0)^T$ in one step, but this is a saddle point.

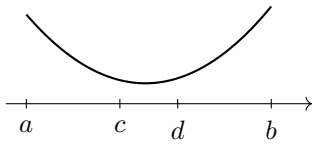
For sensitivity, suppose x^* is a local minimizer with nonsingular Hessian and perturb the objective to $f_b(x) = f(x) - b^T x$. The nearby stationary point satisfies $\nabla f(x_b) = b$, so

$$x_b - x^* \approx H_f(x^*)^{-1} b.$$

Thus the local absolute sensitivity to this linear perturbation is $\left\|H_f(x^*)^{-1}\right\|_2$; small positive Hessian eigenvalues indicate sensitive directions.

Golden section search

(Interval reduction for optimization.) We want to minimize a continuous $f : \mathbb{R} \to \mathbb{R}$ on $[a, b]$. Assume f is *unimodal* on this interval: strictly decreasing up to its minimizer and strictly increasing afterward.



Evaluate f at two interior points $c < d$:

- if $f(c) > f(d)$, the minimizer cannot lie in $[a, c]$—else f would be strictly increasing from c to d, giving $f(c) < f(d)$; retain $[c, b]$;
- if $f(c) \leq f(d)$, retain $[a, d]$. In the equality case, the minimizer lies between c and d, so this is also safe.

To retain the same relative point locations while reusing one function value after each reduction, choose

$$\alpha = \frac{\sqrt{5} - 1}{2} \approx 0.618 : \qquad c = b - \alpha\,(b - a), \qquad d = a + \alpha\,(b - a).$$

Since $\alpha^2 = 1 - \alpha$, one of the previous interior points has the required relative position in the new interval (d becomes the new c, or c the new d). Reusing its function value means that each iteration needs only *one* new function evaluation (two on the first).

Convergence: the bracket length shrinks by the factor α each iteration—linear convergence with constant

$$C = \frac{\sqrt{5} - 1}{2} \approx 0.618,$$

per iteration, and hence also per function evaluation. (Note the contrast: $\varphi = (1 + \sqrt{5})/2 \approx 1.618$ is the convergence *order* of the secant method; golden section search is linear, of order 1, with constant $1/\varphi$.)

Stopping criterion: when $(b - a)/2 \leq \tau$, return the midpoint of the retained interval. Its absolute error in the minimizer's location is at most τ.

CS 450: Numerical Analysis – Lecture 19

Unconstrained Optimization in One and Many Dimensions

Edgar Solomonik

November 3, 2026

Today

Goal: properties and algorithms for unconstrained optimization.

- 1D optimization: conditioning and algorithms
- Steepest descent and conjugate gradient

1D optimization

Given $f : \mathbb{R} \to \mathbb{R}$, solve

$$\min_x f(x).$$

The solution is a *global minimum*; it may not exist if the function diverges to $-\infty$, and it may not be unique. Even a function bounded below need not attain its minimum (e.g. e^x on $\mathbb{R}$). Weaker criterion: a *local minimum*,

$$f(x^*) \leq f(x^* + \delta x) \quad \text{for all sufficiently small } \delta x.$$

When is a local minimum a global minimum?

- hard to ascertain in general;
- true for convex f.

Convexity. A set $S \subset \mathbb{R}^n$ is *convex* if any segment between two points in S is still in S:



convex non-convex

A differentiable function f on an interval is *convex* if its graph lies on or above every tangent line throughout the interval:



convex non-convex

Convex optimization: $\min_{x \in S} f(x)$ with S, f convex. Any local minimum is global. (For twice-differentiable f, convexity on an interval is characterized by $f''(x) \geq 0$ throughout that interval; in n dimensions, the Hessian must be positive semidefinite throughout the convex domain.)

Conditioning of minima. Suppose $|\hat{f}(x) - f(x)| \le \varepsilon$ throughout a neighborhood of a local minimizer x^*, with f twice continuously differentiable and $f''(x^*) > 0$. If $\hat{x}$ minimizes $\hat{f}$ over the same neighborhood and is close to x^*, then

$$f(\hat{x}) \le \hat{f}(\hat{x}) + \varepsilon \le \hat{f}(x^*) + \varepsilon \le f(x^*) + 2\varepsilon.$$

Taylor's expansion at x^*, using $f'(x^*) = 0$, gives

$$f(\hat{x}) - f(x^*) = \tfrac{1}{2} f''(x^*)(\hat{x} - x^*)^2 + \text{h.o.t.}$$

Combining these two relations, the quadratic term can be as large as 2ε, up to higher-order terms. Thus the difference between the minimizers can be as large as approximately

$$|\hat{x} - x^*| \approx 2\sqrt{\frac{\varepsilon}{f''(x^*)}}. \tag{19.1}$$

For small ε, $O(\varepsilon)$ errors in function values can change the minimizer by a much larger amount, $O(\sqrt{\varepsilon})$.

1D algorithms

Golden section search: see last lecture.

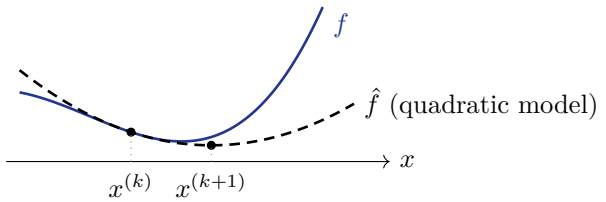
Newton's method. Given $x^{(k)}$, approximate f as

$$f(x) \approx \underbrace{f(x^{(k)}) + f'(x^{(k)})(x - x^{(k)}) + \tfrac{1}{2} f''(x^{(k)})(x - x^{(k)})^2}_{\hat{f}(x)},$$

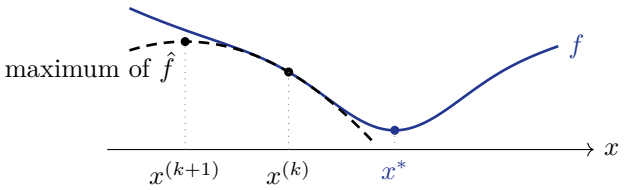
and find a critical point of $\hat{f}$ to obtain $x^{(k+1)}$ (solve $\hat{f}'(x) = 0$); it is the model's minimum when $f''(x^{(k)}) > 0$:

$$\hat{f}'(x) = f'(x^{(k)}) + f''(x^{(k)})(x - x^{(k)}) = 0,$$

$$x^{(k+1)} = x^{(k)} - \frac{f'(x^{(k)})}{f''(x^{(k)})} \qquad \text{(same as a Newton solve on } f'(x) = 0\text{)}. \tag{19.2}$$



Newton's method searches for a *critical point*: if $f''(x^{(k)}) < 0$, the step moves to the *maximum* of $\hat{f}$:



Safeguarding: check $f''(x^{(k)})$ and do the Newton step only if $f''(x^{(k)}) > \varepsilon$ for some small ε. Even then, the full step need not decrease f: shorten it if necessary to obtain decrease, or use a safeguarded alternative such as golden section search on a bracket.

Successive parabolic interpolation. Interpolate $(x^{(k-2)}, f(x^{(k-2)}))$, $(x^{(k-1)}, f(x^{(k-1)}))$, $(x^{(k)}, f(x^{(k)}))$ with a quadratic and use its minimum as the next iterate. If the quadratic has no minimum or proposes a point outside the useful interval, fall back to an interval-reduction step. This is not the secant method for $f'(x) = 0$—that would use $f'(x^{(k-1)})$, $f'(x^{(k)})$—but the convergence is also superlinear, of order $r \approx 1.324$, locally near a sufficiently smooth minimum with positive second derivative.

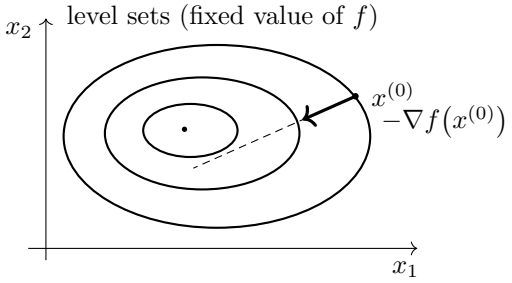
Multidimensional optimization

Given $f : \mathbb{R}^n \to \mathbb{R}$, minimize

$$\min_{x \in \mathbb{R}^n} f(x).$$

Sensitivity to perturbations of the *gradient* depends locally on $\left\| H_f(x^*)^{-1} \right\|_2$, when this inverse exists. This differs from perturbing only function values as above. Algorithms: define in general, analyze for quadratic problems (QPs).

Steepest descent. From $x^{(0)}$, step along the negative gradient:



Solve

$$\min_{\alpha \geq 0} f(x^{(0)} - \alpha \nabla f(x^{(0)}))$$

(a 1-D optimization problem; use e.g. golden-section search or Newton), and repeat from the new iterate. At a noncritical point, a sufficiently small positive step decreases f. Among directions of unit Euclidean norm, the negative-gradient direction gives the greatest first-order decrease.

Convergence analysis of steepest descent. For a QP,

$$f(x) = \tfrac{1}{2} x^T A x - b^T x, \qquad A \text{ SPD}$$

(for a local model in the step variable, use the Hessian as A and the negative gradient as b). The minimizer satisfies $Ax^* = b$. Define the *A-norm*,

$$\|y\|_A = \sqrt{y^T A y},$$

which satisfies the norm properties for SPD A, and

$$f(x) - f(x^*) = \tfrac{1}{2} \|x - x^*\|_A^2 .$$

With exact line searches, a worst-case linear convergence factor in the A-norm is

$$\frac{\kappa_2(A) - 1}{\kappa_2(A) + 1}. \tag{19.3}$$

Actual convergence can be faster; a large condition number makes this worst-case factor close to one.

Conjugate gradient (linear CG). Gradient descent with extrapolation:

$$\min_{\alpha,\beta} \; f\Big(x^{(k)} - \alpha \nabla f\big(x^{(k)}\big) + \beta \underbrace{\big(x^{(k)} - x^{(k-1)}\big)}_{\text{“momentum”}}\Big).$$

For general f, the parallel-tangents method (PARTAN) is a nonlinear CG variant: take a steepest-descent step, then minimize along the line through that new point and $x^{(k-1)}$. For an SPD QP with exact line searches, it reproduces linear CG, which solves the two-parameter problem above (after an initial steepest-descent step). The convergence statements below concern the SPD quadratic case. Its worst-case convergence bound is governed by the linear factor

$$\frac{\sqrt{\kappa_2(A)} - 1}{\sqrt{\kappa_2(A)} + 1}. \tag{19.4}$$

This describes progress over multiple iterations, not a fixed error ratio at each step. CG searches A-orthogonal directions, $s^{(i)T} A s^{(j)} = 0$ for $i \neq j$, and converges in at most n iterations in exact arithmetic. Rounding errors can prevent exact termination in n steps.

Equivalently, starting with $x^{(0)} = 0$, the kth CG iterate minimizes $f(x)$ for

$$x \in \mathrm{Kr}_k(b) = \mathrm{span}\{b,\, Ab,\, \ldots,\, A^{k-1}b\},$$

i.e., CG is Krylov optimization for QPs; the space is seeded by $b = r^{(0)}$, the residual at $x^{(0)} = 0$, for the QP $f(x) = \frac{1}{2} x^T A x - x^T b$. With Q_k an orthonormal basis of $\mathrm{Kr}_k(b)$,

$$\min_{x \in \mathrm{Kr}_k(b)} f(x) = \min_{y \in \mathbb{R}^k} f(Q_k y) = \min_{y} \; \tfrac{1}{2}\, y^T \underbrace{Q_k^T A Q_k}_{H_k}\, y - y^T Q_k^T b,$$

so $y^{(k)} = H_k^{-1} Q_k^T b$ and $x^{(k)} = Q_k y^{(k)}$. H_k changes by the addition of one row and one column at each step. After padding H_k with an extra diagonal entry, this is a low-rank update, so Sherman–Morrison gives intuition for an efficient update from the previous reduced solution to the next. With a Lanczos basis, H_k is tridiagonal. CG uses this structure to update the approximation, residual, and search direction with one matrix-vector product and $O(n)$ additional work, as in the formulas below.

Full CG algorithm for QPs. (Or for solving $Ax = b$ with SPD A.)

- $x^{(0)} = 0; \quad r^{(0)} = s^{(0)} = b - A x^{(0)}$;
- for $k = 0, 1, \ldots$ while $\left\| r^{(k)} \right\|_2 > \tau$, with a prescribed absolute tolerance $\tau > 0$:

$$\alpha = \frac{r^{(k)T} r^{(k)}}{s^{(k)T} A s^{(k)}} \qquad \text{(solve for the step size in the direction of } s^{(k)}\text{)}$$
$$x^{(k+1)} = x^{(k)} + \alpha\, s^{(k)}$$
$$r^{(k+1)} = r^{(k)} - \alpha\, A s^{(k)} \qquad \text{(update the residual } r^{(k+1)} = b - A x^{(k+1)} = -\nabla f(x^{(k+1)})\text{)}$$
$$\beta = \frac{r^{(k+1)T} r^{(k+1)}}{r^{(k)T} r^{(k)}} \qquad \text{(solve for the momentum parameter)}$$
$$s^{(k+1)} = r^{(k+1)} + \beta\, s^{(k)} \qquad \text{(define the new search direction)}$$

Each step uses one matrix-vector product and $O(n)$ vector work, with a fixed number of length-n working vectors; no full Krylov basis need be stored.

CS 450: Numerical Analysis – Lecture 20

Quasi-Newton Methods, Nonlinear Least Squares, and Constrained Optimization

Edgar Solomonik

November 5, 2026

Today

Goal: go beyond unconstrained nonlinear and quadratic optimization.

- Quasi-Newton methods
- Nonlinear least squares
- Constrained optimization: properties of problem and solutions; SQP and active-set methods

Review: Newton and quasi-Newton methods

Newton's method in n dimensions is the fixed-point iteration on

$$g(x) = x - \underbrace{H_f(x)^{-1}}_{\substack{\text{expensive to} \\ \text{compute and use}}} \nabla f(x).$$

Local quadratic convergence holds near a sufficiently smooth minimum with positive definite Hessian, but steps are expensive and can be unreliable farther away; compute the step by solving a linear system, not forming the inverse. A *line search* chooses a step length $\alpha > 0$ along a given direction p by considering the one-dimensional function $f(x^{(k)} + \alpha p)$, for example by approximately minimizing it or finding a step that decreases it.

Quasi-Newton methods. Replace the Hessian by an approximation $\hat{H}_k \approx H_f(x^{(k)})$:

$$x^{(k+1)} = x^{(k)} - \hat{H}_k^{-1} \nabla f(x^{(k)}).$$

With

$$y_k = \nabla f(x^{(k)}) - \nabla f(x^{(k-1)}), \qquad s_k = x^{(k)} - x^{(k-1)},$$

we impose the *secant condition*

$$\hat{H}_k \, s_k = y_k;$$

we also want $\hat{H}_k$ SPD.

BFGS algorithm. Start from $\hat{H}_0 = I$ (or $\hat{H}_0 = H_f(x^{(0)})$ if it is positive definite). Assume the step lengths and updates are chosen so that the Hessian approximations remain positive definite. The BFGS update is

$$\hat{H}_k = \hat{H}_{k-1} + \frac{y_k \, y_k^T}{y_k^T s_k} - \frac{\hat{H}_{k-1} \, s_k \, s_k^T \hat{H}_{k-1}}{s_k^T \hat{H}_{k-1} \, s_k} \qquad \text{(a correction of rank at most 2).} \tag{20.1}$$

- solves the secant condition ($\hat{H}_k \, s_k = y_k$);
- preserves the symmetry expected from H_f;
- Sherman–Morrison–Woodbury $\Rightarrow O(n^2)$ cost per step to update the inverse or a factorization (preferred for stability), rather than refactoring from scratch; this excludes the cost of evaluating f and ∇f;
- no second-order derivatives needed.

At a noncritical point, a positive definite $\hat{H}_k$ gives a *descent direction* $p = -\hat{H}_k^{-1} \nabla f(x^{(k)})$: a sufficiently small positive step length α gives $f(x^{(k)} + \alpha p) < f(x^{(k)})$. Use a line search to choose α.

Nonlinear least squares

Motivation: fitting a nonlinear model. Given measurements (t_i, d_i), fit a model $\phi(t; x)$ with unknown parameters $x \in \mathbb{R}^n$ so that $\phi(t_i; x) \approx d_i$. Define

$$r_i(x) = d_i - \phi(t_i; x), \qquad \min_x f(x) = \frac{1}{2}\sum_{i=1}^{m} r_i(x)^2 = \frac{1}{2}\,\|r(x)\|_2^2 .$$

Often there are more measurements than parameters ($m > n$), and noise or model error can prevent all residuals from being zero. Least squares seeks parameters that fit the measurements as closely as possible in the sum-of-squares sense. If the model is affine in the unknown parameters, this is linear least squares. Otherwise, it is nonlinear least squares. It is the dependence on the parameters, not on t, that matters.

Example: a decay curve with an unknown baseline. Fit a signal approaching a baseline by

$$\phi(t; x) = x_1 e^{-x_2 t} + x_3,$$

where x_1 is the initial amount above the baseline, x_2 is the decay rate, and x_3 is the baseline. For example, consider the illustrative measurements

t_i	0	1	2	3
d_i	5.0	3.0	2.2	1.6

and minimize $\frac{1}{2}\sum_{i=1}^{4}(d_i - x_1 e^{-x_2 t_i} - x_3)^2$. There are four measurements and three parameters. The residuals are nonlinear in x_2; their Jacobian has rows

$$[J_r(x)]_{i,:} = \begin{bmatrix} -e^{-x_2 t_i} & x_1 t_i e^{-x_2 t_i} & -1 \end{bmatrix}.$$

Gauss–Newton uses these derivatives to compute a parameter correction by linear least squares.

Gauss–Newton. For $f(x) = \frac{1}{2}\,\|r(x)\|_2^2$,

$$\nabla f(x) = J_r(x)^T r(x), \qquad H_f(x) = J_r(x)^T J_r(x) + \underbrace{\sum_{i=1}^{m} r_i(x)\, H_{r_i}(x)}_{\substack{\text{small for small residuals and} \\ \text{bounded residual Hessians}}} .$$

Gauss–Newton algorithm: assuming $J_r(x)$ has full column rank, drop the second term,

$$\hat{H}_f(x) = J_r(x)^T J_r(x),$$

$$x^{(k+1)} = x^{(k)} - \hat{H}_f\big(x^{(k)}\big)^{-1} \nabla f\big(x^{(k)}\big), \tag{20.2}$$

Linearizing the residual gives

$$r\big(x^{(k)} + s\big) \approx r\big(x^{(k)}\big) + J_r\big(x^{(k)}\big) s.$$

Thus the step is computed as the linear least-squares problem

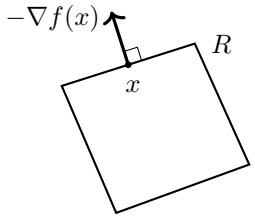
$$\min_s \big\| J_r\big(x^{(k)}\big)\, s + r\big(x^{(k)}\big) \big\|_2,$$

followed by $x^{(k+1)} = x^{(k)} + s$; solve the linear least-squares problem by QR rather than explicitly forming and solving the normal equations. Thus nonlinear least squares is solved by a sequence of linear least-squares problems. If r is twice continuously differentiable near x^*, $r(x^*) = 0$, and $J_r(x^*)$ has full column rank, Gauss–Newton converges locally at least quadratically. A least-squares minimizer need not have zero residual; then quadratic convergence is not guaranteed, and a large residual can cause slow convergence or failure. The regularized version is called Levenberg–Marquardt.

Constrained optimization

$$\min_x f(x) \quad \text{s.t.} \quad g(x) = 0, \quad h(x) \le 0 \qquad \Longleftrightarrow \qquad \min_{x \in R} f(x), \quad \underbrace{R = \{x : g(x) = 0,\ h(x) \le 0\}}_{\text{feasible region}}.$$

When is x a local minimum. At a feasible point x, $\nabla f(x) = 0$ and a positive definite $H_f(x)$ are sufficient for a strict local minimum, but neither is necessary. Assume the objective and constraints are twice continuously differentiable.



For equality constraints, assume that the constraint gradients are linearly independent at x. A direction tangent to the constraints satisfies $J_g(x)s = 0$. At a local minimum, $\nabla f(x)^T s = 0$ for all such directions, so

$$\nabla f(x) \in \operatorname{span}(\{\nabla g_1(x), \ldots, \nabla g_m(x)\}),$$

where m is the number of equality constraints. For nonlinear constraints, a straight step $x + \epsilon s$ along a tangent direction need not remain feasible.

Full necessary conditions: Karush–Kuhn–Tucker (KKT). Assume that the gradients of the equality constraints and active inequality constraints are linearly independent at the solution. Then a local minimum satisfies the following conditions.

- x is feasible: $g(x) = 0$, $h(x) \le 0$.
- Let the *active* inequality constraints be

  $$\hat{h}(x) = \begin{bmatrix} h_{i_1}(x) \\ \vdots \\ h_{i_k}(x) \end{bmatrix},$$

  where $h_{i_1}, \ldots, h_{i_k}$ are the components of h that hold with equality, $h_{i_1}(x) = \cdots = h_{i_k}(x) = 0$, hence $\hat{h}(x) = 0$. An inactive inequality satisfies $h_i(x) < 0$.
- There exist λ_1, λ_2 such that

  $$-\nabla f(x) = J_g(x)^T \lambda_1 + J_h(x)^T \lambda_2, \qquad \lambda_2 \ge 0. \tag{20.3}$$

  The equality multipliers λ_1 have no sign restriction. With $\lambda = (\lambda_1, \lambda_2)$, x is stationary with respect to x for the *Lagrangian function*

  $$\mathcal{L}(x, \lambda) = f(x) + \lambda_1^T g(x) + \lambda_2^T h(x).$$
- *Complementary slackness:* $(\lambda_2)_i h_i(x) = 0$ for each inequality. An inactive inequality has zero multiplier; an active inequality may have a nonzero multiplier.

For the single boundary pictured below, the sign condition $\lambda_2 \ge 0$ distinguishes an outward-pointing $-\nabla f$ from an inward-pointing one that allows f to decrease inside:



The KKT conditions alone do not establish a local minimum. Second-order conditions concern the Hessian of the Lagrangian with respect to x (holding the multipliers fixed), denoted $B(x, \lambda)$, and its behavior along directions allowed by the constraints. For nonlinear constraints, H_f alone does not suffice; for linear constraints, $B = H_f$.

Algorithms for constrained optimization

The first-order KKT conditions lead to nonlinear equations for stationarity and feasibility, together with sign and complementarity conditions when inequalities are present.

Sequential quadratic programming (SQP). Assume equality constraints only. The local equations in the variables (x, λ) are

$$g(x) = 0, \qquad J_g(x)^T \lambda = -\nabla f(x),$$

i.e.,

$$F(x, \lambda) = \begin{bmatrix} \nabla f(x) + J_g(x)^T \lambda \\ g(x) \end{bmatrix} = 0. \tag{20.4}$$

Apply Newton's method to these stationarity equations, not to minimizing $\mathcal{L}$ jointly over x and λ. Solve

$$J_F\big(x^{(k)}, \lambda^{(k)}\big) \begin{bmatrix} s \\ \delta\lambda \end{bmatrix} = -F\big(x^{(k)}, \lambda^{(k)}\big), \qquad \begin{aligned} x^{(k+1)} &= x^{(k)} + s, \\ \lambda^{(k+1)} &= \lambda^{(k)} + \delta\lambda, \end{aligned}$$

where

$$J_F(x, \lambda) = \begin{bmatrix} H_f(x) + \sum_i \lambda_i \overbrace{H_{g_i}(x)}^{0 \text{ if } g \text{ is linear}} & J_g(x)^T \\ J_g(x) & 0 \end{bmatrix}.$$

The upper-left block is $B(x, \lambda)$. This symmetric block matrix is generally indefinite, so ordinary Cholesky and the positive-definite version of CG do not apply directly.

The step is also obtained from the stationarity equations of a quadratic subproblem with linearized constraints:

$$\min_s \ \nabla f\big(x^{(k)}\big)^T s + \tfrac{1}{2} s^T B_k s \quad \text{s.t.} \quad g\big(x^{(k)}\big) + J_g\big(x^{(k)}\big) s = 0,$$

where $B_k = B\big(x^{(k)}, \lambda^{(k)}\big)$ or an approximation to it. With the exact Hessian, the subproblem's multiplier is $\lambda^{(k)} + \delta\lambda$ in the Newton system above. The Hessian approximation may need modification so that the quadratic subproblem has a minimizer.

- A fixed point of SQP satisfies the first-order KKT equations; it need not be a local minimum.
- With exact Hessians, sufficient smoothness, and a nonsingular J_F at the solution, local quadratic convergence is possible from sufficiently close starting points.
- Quasi-Newton approximations to the Lagrangian Hessian may be preferable in practice.

Active-set methods for inequality constraints.

- let $\hat{h}_k$ collect the inequality constraints in the current working set. For feasible active-set iterates, $\hat{h}_k(x^{(k)}) = 0$;
- treat $\hat{h}_k(x) = 0$, together with $g(x) = 0$, as equality constraints and take a step of the corresponding subproblem, for example an SQP step;
- add constraints that block the proposed step; at a solution of the working-set subproblem, remove constraints whose multipliers have the wrong sign (negative under our convention);
- check feasibility and progress in the objective. For linear constraints, shorten the step as needed to remain feasible. For nonlinear constraints, satisfying the linearized constraints, even with a shorter step, need not preserve feasibility for the original problem.

CS 450: Numerical Analysis – Lecture 21

Review of Nonlinear Optimization; Penalty and Barrier Methods

Edgar Solomonik

November 10, 2026

Today

Goal: review nonlinear optimization with a focus on advanced topics.

- n-dimensional optimization methods
- Nonlinear least squares
- Constrained optimization

n-dimensional optimization algorithms

$$\min_{x\in\mathbb{R}^n} f(x).$$

Steepest descent.

$$x^{(k+1)} = x^{(k)} - \alpha\nabla f(x^{(k)}), \qquad \alpha \text{ the solution to the 1-D optimization problem } \min_{\alpha\geq 0} f(x^{(k)} - \alpha\nabla f(x^{(k)})).$$

If $f(x) = \frac{1}{2}x^TAx - x^Tb$ with SPD A, exact line searches give linear convergence in the A-norm with worst-case factor

$$C = \frac{\kappa_2(A)-1}{\kappa_2(A)+1}$$

For general smooth f, a quadratic model near a minimum with positive definite Hessian gives the corresponding local picture. Consecutive search directions are orthogonal under an exact line search:

$$\frac{\partial}{\partial\alpha} f(x^{(k)} - \alpha\nabla f(x^{(k)})) = 0 \quad \Longrightarrow \quad \nabla f(\underbrace{x^{(k)} - \alpha\nabla f(x^{(k)})}_{x^{(k+1)}})^T\nabla f(x^{(k)}) = 0.$$

Conjugate gradient. Nonlinear version (of the linear CG from Lecture 19):

- search direction: with $s^{(0)} = -\nabla f(x^{(0)})$, for $k \geq 1$,

  $$s^{(k)} = -\nabla f(x^{(k)}) + \beta\, s^{(k-1)}, \qquad x^{(k+1)} = x^{(k)} + \alpha\, s^{(k)},$$

  with $\alpha > 0$ chosen by a line search along $s^{(k)}$ (also at $k = 0$);
- β is computed by a formula such as Fletcher–Reeves,

  $$\beta = \frac{\nabla f(x^{(k)})^T\nabla f(x^{(k)})}{\nabla f(x^{(k-1)})^T\nabla f(x^{(k-1)})}.$$

For the SPD quadratic $f(x) = \frac{1}{2}x^TAx - x^Tb$, with exact line searches and exact arithmetic:

- the usual CG direction formulas agree, and the new iterate minimizes

  $$f(x^{(k)} + \gamma_1\nabla f(x^{(k)}) + \gamma_2 s^{(k-1)}) \quad \text{over } \gamma_1 \text{ and } \gamma_2$$

  (the two-parameter problem from Lecture 19, for $k \geq 1$);
- the standard worst-case convergence factor is

  $$\frac{\sqrt{\kappa_2(A)}-1}{\sqrt{\kappa_2(A)}+1};$$

  it describes the overall reduction of the A-norm error, not necessarily the reduction at each individual step;

- *all* search directions are A-orthogonal: $s^{(k)T} A s^{(j)} = 0$ for $j \neq k$;
- convergence in at most n steps;
- CG solves $Ax = b$, since $\nabla f(x) = 0$ means $Ax = b$.

The conjugacy and finite-termination statements do not extend to general nonlinear f. Apart from the line search, steepest descent and nonlinear CG require one new gradient evaluation and $O(n)$ vector work per iteration; the line search may need additional evaluations.

Newton's method.

$$H_f(x^{(k)}) s^{(k)} = -\nabla f(x^{(k)}), \qquad x^{(k+1)} = x^{(k)} + s^{(k)}.$$

With sufficient smoothness, convergence is locally at least quadratic near a critical point with nonsingular Hessian. The critical point need not be a minimum, and the steps are expensive.

BFGS.

- $\hat{H}_k s^{(k)} = -\nabla f(x^{(k)})$ with $x^{(k+1)} = x^{(k)} + s^{(k)}$;
- assume the step lengths and updates are chosen so that $\hat{H}_k$ remains SPD;
- the update enforces the secant equation

$$\hat{H}_k(x^{(k)} - x^{(k-1)}) = \nabla f(x^{(k)}) - \nabla f(x^{(k-1)});$$

- $\operatorname{rank}(\hat{H}_k - \hat{H}_{k-1}) \leq 2$: Sherman–Morrison–Woodbury motivates updating the inverse, or a factorization, for $O(n^2)$ linear-algebra work per step with dense $\hat{H}_k$, after initialization;
- at a noncritical point, $s^{(k)}$ is a descent direction: a sufficiently small positive α gives

$$f(x^{(k)} + \alpha s^{(k)}) < f(x^{(k)}).$$

Lemma. For any SPD matrix A, the direction $p = -A\nabla f(x)$ is a descent direction from x: wherever $\nabla f(x) \neq 0$,

$$\frac{\partial}{\partial \alpha} f(x + \alpha p) = \nabla f(x + \alpha p)^T p,$$

and at $\alpha = 0$ we have

$$\frac{\partial}{\partial \alpha} f(x + \alpha p)\Big|_{\alpha=0} = \nabla f(x)^T p = -\nabla f(x)^T A \nabla f(x) < 0.$$

This applies to the BFGS step $s^{(k)} = -\hat{H}_k^{-1} \nabla f(x^{(k)})$, i.e. $A = \hat{H}_k^{-1}$, which is SPD whenever $\hat{H}_k$ is.

Nonlinear least squares

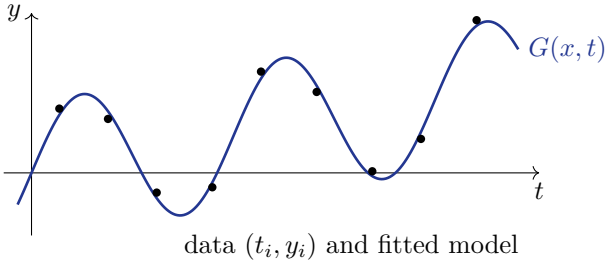
Given data

$$(t_1, y_1), \ \ldots, \ (t_m, y_m),$$

we want a model $G(x, t)$ with parameters $x \in \mathbb{R}^n$ solving

$$\min_x \ f(x), \qquad f(x) = \frac{1}{2} \sum_{i=1}^{m} \big(G(x, t_i) - y_i\big)^2, \qquad \text{e.g.} \quad G(x, t) = x_1 t + \sin(2\pi x_2 t).$$

The coefficient x_1 controls the linear trend; x_2 controls the oscillation frequency. This model is nonlinear in x_2 and is defined for all $x_1, x_2 \in \mathbb{R}$. Noise or model error may prevent an exact fit to all measurements.



data (t_i, y_i) and fitted model

Gauss–Newton algorithm. At the iterate $x^{(k)}$, approximate

$$u(x) = \begin{bmatrix} G(x, t_1) \\ \vdots \\ G(x, t_m) \end{bmatrix} \qquad \text{linearly:} \qquad \hat{u}(x) = u(x^{(k)}) + J_u(x^{(k)})(x - x^{(k)}).$$

With $y = (y_1, \ldots, y_m)^T$ and the step $s = x - x^{(k)}$, minimize

$$\hat{f}(x) = \frac{1}{2}\|\hat{u}(x) - y\|_2^2 = \frac{1}{2}\|\underbrace{J_u(x^{(k)})}_{A}\, s - \underbrace{(y - u(x^{(k)}))}_{b}\|_2^2,$$

a *linear least-squares problem* $\min_s \|As - b\|_2$; then set $x^{(k+1)} = x^{(k)} + s$. In last lecture's notation the residual is $r(x) = y - u(x)$, so $f(x) = \frac{1}{2}\|r(x)\|_2^2$ and this is the same Gauss–Newton step. Here J_u is $m \times n$; the linear least-squares problem can be solved by QR.

A least-squares solution need not have $r(x^*) = 0$. If r is twice continuously differentiable near x^*, $r(x^*) = 0$, and $J_r(x^*)$ has full column rank, Gauss–Newton converges locally at least quadratically. With nonzero residual, this rate is not generally guaranteed; convergence may be slower or may fail.

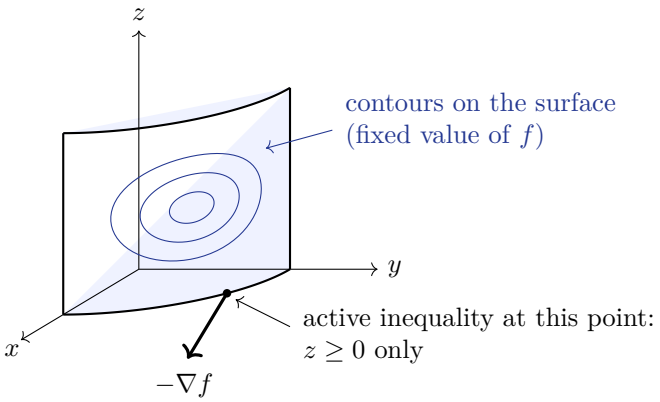
Constrained optimization

$$\min_{x \in \mathbb{R}^n} f(x), \qquad g(x) = 0, \quad h(x) \le 0.$$

Example.

$$\min_{x,y,z} f(x, y, z), \qquad x^2 + y^2 = 1, \quad x, y, z \ge 0.$$

The feasible set is the quarter-cylinder surface over the arc of the unit circle in the first octant:



When is (x, y, z) a local minimum? Assuming the equality and active inequality gradients are linearly independent, necessary first-order conditions are:

- It must satisfy the constraints.
- The negative gradient

$$-\begin{bmatrix} \partial f/\partial x \\ \partial f/\partial y \\ \partial f/\partial z \end{bmatrix}(x, y, z)$$

must be a linear combination of the equality-constraint gradients and active inequality-constraint gradients.

At the point pictured, $x > 0$, $y > 0$, and $z = 0$, so only $z \ge 0$ is active. The equality and active inequality are

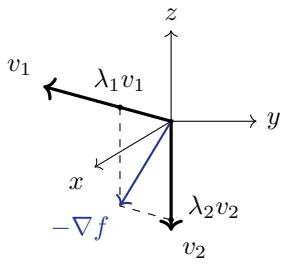
$$g(x, y, z) = 1 - x^2 - y^2 = 0, \qquad \hat{h}(x, y, z) = -z \quad \text{(the active inequality } z \ge 0\text{, written as } -z \le 0\text{)},$$

with

$$J_{[g;\hat{h}]}(x,y,z) = \begin{bmatrix} \nabla g(x,y,z)^T \\ \nabla \hat{h}(x,y,z)^T \end{bmatrix} = \begin{bmatrix} -2x & -2y & 0 \\ 0 & 0 & -1 \end{bmatrix} = \begin{bmatrix} v_1(x,y,z)^T \\ v_2(x,y,z)^T \end{bmatrix}.$$

If this point is a local minimum, there exist λ_1, λ_2 such that

$$-\nabla f(x,y,z) = \lambda_1 v_1(x,y,z) + \lambda_2 v_2(x,y,z). \tag{21.1}$$



Here λ_1 is unrestricted and $\lambda_2 \geq 0$, as in last lecture's KKT conditions. The third component gives $\partial f/\partial z = \lambda_2 \geq 0$: moving toward positive z cannot decrease f to first order at a minimum. At $x = 0$ or $y = 0$, include the corresponding nonnegativity constraint among the active constraints. These first-order conditions alone do not establish a local minimum.

Algorithms for constrained optimization

Solve the first-order KKT conditions. Under the constraint-gradient independence assumption, a constrained local minimum satisfies these conditions; satisfying them alone does not establish a minimum. For equality constraints only, the Lagrangian

$$\mathcal{L}(x,\lambda) = f(x) + g(x)^T\lambda$$

must have a critical point at a constrained minimum, so

$$\nabla\mathcal{L}(x,\lambda) = 0.$$

If f is quadratic with symmetric A and g is linear,

$$f(x) = \tfrac{1}{2}x^TAx - b^Tx, \qquad g(x) = Bx,$$

we have the *saddle-point system*

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} x \\ \lambda \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}, \tag{21.2}$$

where $A = H_f(x)$ and $B = J_g(x)$. The system matrix is symmetric but, with constraints present, is not positive definite even when A is positive definite. Thus the SPD version of CG does not apply directly to this full system.

Penalty and barrier functions

Modify f to obtain a sequence of optimization problems whose solutions approach a constrained solution under suitable assumptions.

Penalty (for equality constraints g).

$$\min_x f(x) \text{ s.t. } g(x) = 0 \qquad \Longrightarrow \qquad \min_x f(x) + \frac{\rho}{2}\, g(x)^T g(x), \tag{21.3}$$

solved for a sequence of increasing $\rho > 0$ (equivalently, $\rho = 1/\mu$ with decreasing $\mu > 0$). The penalty adds a cost for violating the equalities; intermediate iterates need not be feasible. A finite penalty parameter does not generally enforce the equalities exactly.

Barrier functions. (Interior point methods, for inequality constraints.)

$$\min_x f(x) \text{ s.t. } h(x) \leq 0 \qquad \Longrightarrow \qquad \min_x f(x) - \mu \log(-h(x)) \quad \text{or} \quad \min_x f(x) - \mu \frac{1}{h(x)}, \tag{21.4}$$

with $\mu > 0$ decreasing toward zero. These subproblems are restricted to $h(x) < 0$; for several inequalities, sum the barrier terms over all components, e.g. $-\mu \sum_i \log(-h_i(x))$. Both added terms tend to $+\infty$ as $h(x) \to 0^-$. The logarithm is not defined on or outside the boundary, and the reciprocal expression must not be used outside the feasible interior.

The basic barrier method starts strictly inside the inequalities and keeps accepted iterates there. Equality constraints, if present, are handled separately. As ρ increases for a penalty method, or μ decreases for a barrier method, the subproblems can become increasingly ill-conditioned. Change the parameter gradually, using each approximate solution to initialize the next subproblem.

The logarithmic barrier is the most common choice. For an inequality-only problem, its connection to the KKT conditions is obtained by setting $\lambda_i = -\mu/h_i(x) > 0$. At a stationary point of the barrier objective,

$$\nabla f(x) + J_h(x)^T \lambda = 0, \qquad \lambda_i h_i(x) = -\mu.$$

Thus the complementary products approach zero as $\mu \to 0$.

AI-use disclosure. *These notes were prepared with the assistance of Claude and ChatGPT, including transcription from handwritten lecture notes and drafting or editing of prose, examples, and derivations. The author reviewed the final document and assumes responsibility for its content.*