Today:

- Trust-Region Methods (Ch. 7, Ch. 8)
- Global Optimization with Simulated Annealing (Ch. 18)

Plan for next week:

- Tuesday online lecture (face to face option also)
- Thursday:
  - wrap-up
  - structure of Exam #1
  - Coding session (bring laptops)

Quick recap: Trust-region methods. Idea is to approximate $f(\underline{x}_k+\underline{p})$ as a QMP:

$$\left.\begin{aligned} f(\underline{x}_k+\underline{p}) &\simeq f_k + \langle \underline{g},\underline{p}\rangle + \tfrac{1}{2}\langle \underline{p}, B\underline{p}\rangle \\ &= m_k(\underline{p}) \end{aligned}\right\} \quad (1)$$

This is a valid approximation so long as $\|\underline{p}\|_2$ is not too large, i.e. provided we are in the trust region defined by:

$$\|\underline{p}\|_2 \leq \Delta \; .$$

At each iteration k, we choose a $\underline{p}_*$ that satisfies:

$$\underline{p}_* = \underset{\|\underline{p}\|_2 \leq \Delta}{\arg\min} \; m_k(\underline{p}) \qquad (2)$$

Then, $\underline{x}_{k+1} = \underline{x}_k + \underline{p}_*$.

Theorem 7.1 tells us how to solve for $\underline{p}_*$ exactly.

Theorem 4.1 tells us how to solve for p_* exactly.

Two approximate methods:

- Cauchy Point method ☑
- Dogleg method

Cauchy Point Method Back to QMP. If Δ is sufficiently small, and $\|p\|_2 \le \Delta$, then $m_k(p)$ can be approximated as:

$$m_k(p) \simeq f_k + \underbrace{\langle g, p \rangle}$$

The p that minimizes $m_k(p)$ is then $\propto -g$.

Take:

$$p_{temp} = -\frac{\Delta}{\|g\|_2} g$$

Then take

$$p = \tau p_{temp}, \quad \text{where } \tau > 0 \text{ is TBC.}$$

We find τ by solving:

$$\tau = \arg\min_{\tau > 0} m_k(\tau p_{temp})$$

We compute τ from this formula (EXAM)

$$m_k(\tau p_{temp}) = f_k + \tau \langle g, p_{temp} \rangle + \frac{1}{2}\tau^2 \langle p_{temp}, B p_{temp} \rangle$$

Sub in for

$$p_{temp} = -\frac{\Delta}{\|g\|_2} g$$

Hence:

$$m_k(\tau p_{temp}) = f_k - \tau \frac{\Delta}{\|g\|_2} \overbrace{\langle g, g\rangle}^{=\|g\|_2^2}$$

$$+ \frac{1}{2}\tau^2 \frac{\Delta}{\|g\|_2}\frac{\Delta}{\|g\|_2} \langle g, Bg\rangle$$

$$\Rightarrow m_k(\tau p_{temp}) = \underbrace{f_k - \tau \Delta \|g\|_2 + \frac{1}{2}\tau^2 \frac{\Delta^2}{\|g\|_2^2}\langle g, Bg\rangle}_{Q(\tau)}$$

We find the τ that minimizes $Q(\tau)$.

Case 1: $\boxed{\langle g, Bg\rangle \leq 0}$. Choose τ as large as possible to make $Q(\tau)$ as small as possible. Hence $\tau = 1$ (stay inside trust region), hence

$$p = -\frac{\Delta}{\|g\|_2} g.$$

Case 2: $\langle g, Bg\rangle > 0$

$$\overset{=0}{Q'(\tau)} = -\cancel{\Delta}\|g\|_2 + \tau \frac{\Delta^{\cancel{2}}}{\|g\|_2^2} \underbrace{\langle g, Bg\rangle}_{>0}$$

To find the min, set $Q'(\tau) = 0$, hence:

$$\tau = \frac{\|g\|_2^3}{}$$

$$\Delta \langle g, Bg \rangle$$

To stay inside the trust region, we require $0 \le \tau \le 1$, hence

$$\tau = \min\left(1, \frac{\|g\|_2^3}{\Delta \langle g, Bg \rangle}\right)$$

Summarizing,

$$p_{cauchy} = -\tau \frac{\Delta}{\|g\|_2} g$$

where

$$\tau = \begin{cases} 1 & \text{if } \langle g, Bg \rangle \le 0 \\ \min\left(1, \frac{\|g\|_2^3}{\Delta \langle g, Bg \rangle}\right), & \langle g, Bg \rangle > 0 \end{cases} \quad \blacksquare$$

Observation: We have

$$p_{cauchy} \propto -g$$

The obvious choice for g is ∇f_k, giving

$$p_{cauchy} \propto -\nabla f_k$$

So the Cauchy-Point Method is SD in disguise (Linear Convergence Rate).

For this reason, we look at a better way of finding approximate solutions to (2) - Dogleg Method.

Back to (2) :

$$\text{minimize } m_k(p) = f_k + \langle g, p\rangle + \frac{1}{2}\langle p, Bp\rangle$$

Subject to $\|p\|_2 \le \Delta$.

Try: $\nabla_p m_k(p) = 0$. This gives:

$$Bp = -g \quad ,$$

hence $p = -B^{-1}g$, provided B is invertible.

Check: Is $\|p\|_2 \le \Delta$? If so, then we have an exact solution to Equation (2).

What if $\|p\|_2 >$? Then, we construct an approximate solution. First, look at:

$$p_{SD} = -\frac{\langle g, g\rangle}{\langle g, Bg\rangle} g$$

If p_{SD} is outside the trust region, then we must accept

$$p_* = -\frac{\Delta}{\|g\|_2} g$$

as the approximate solution of (2).

If p_{SD} is inside the trust region, then we can accept

$$p_* = p_{SD} + \alpha(p_{NEWTON} - p_{SD}) \quad , \quad \alpha \in [0,1] \quad .$$

We choose α such that p_* sits on the trust-region boundary.

region boundary.

Summarizing:

- Case 1, if $\|p_{NEWTON}\|_2 \le \Delta$, we accept
$$p_* = p_{NEWTON}$$
as the (exact) solution of (2).

- Case 2, if $\|p_{NEWTON}\|_2 > \Delta$
  - Case 2.1 If $\|p_{SD}\|_2 > \Delta$, then we accept
$$p_* = -\frac{\Delta}{\|g\|_2} g$$
as an approximate solution of (2)
  - Case 2.2 If $\|p_{SD}\|_2 < 0$, then we accept
$$p_* = p_{SD} + \alpha(p_{NEWTON} - p_{SD}), \quad \alpha \in [0,1]$$
as an approximate solution of (2).

Analysis of the Dogleg Method (§8.2)

We want to show: $\exists\ \alpha \in [0,1]$ such that $\|p_*\|_2 = \Delta$.

i.e. we want to show:
$$\|p_*\|_2^2 = \Delta^2$$
i.e. $\|p_{SD} + \alpha(p_{NEWTON} - p_{SD})\|_2^2 = \Delta^2$

i.e. $\|\underline{p}^{SD} + \underline{\alpha}(\underline{p}^N - \underline{p}^{SD})\|_2^2 = \Delta^2$

i.e.

$$\overset{a}{\alpha^2 \|\underline{p}^N - \underline{p}^{SD}\|_2^2} + 2\alpha \overset{b}{\langle \underline{p}^{SD}, \underline{p}^N - \underline{p}^{SD} \rangle} + \overset{c}{\|\underline{p}^{SD}\|_2^2 - \Delta^2} = 0$$

Quadratic in α.

i.e. we want to show

$$a\alpha^2 + 2b\alpha + c = 0$$

$$\alpha = \frac{-b \pm \sqrt{b^2 - ac}}{a}$$

Check first that α is real. So look at the sign of the discriminant $b^2 - ac$. In particular, $c = \|\underline{p}^{SD}\|_2^2 - \Delta^2 < 0$ (Case 2.2), so the discriminant is positive, so α is real.

Also, look at:

$$h(\alpha) = \|\underline{p}^{SD} + \overset{=1}{\alpha}(\underline{p}^N - \underline{p}^{SD})\|_2^2$$

$$h(0) = \|\underline{p}^{SD}\|_2^2 \quad (< \Delta^2)$$

$$h(1) = \|\underline{p}^N\|_2^2 \quad (> \Delta^2)$$

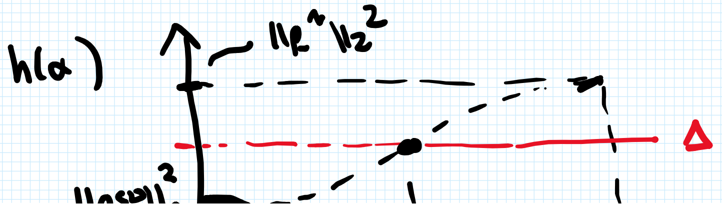

By continuity of $h(\alpha)$ on $[0,1]$, there exists $\alpha_0 \in [0,1]$ such that $h(\alpha_0) = \Delta^2$.

Key property of Dogleg Method: Superlinear Convergence rate:

$$\|\underline{x}_{k+1} - \underline{x}_*\|_2 \leq C \|\underline{x}_k - \underline{x}_*\|_2^{1+\epsilon}, \quad \epsilon > 0.$$

Trust-region methods (with an approximate solution of the sub-problem in Equation (2)) can be used in cases when B is either P.D. or not P.D.

Three cases:

1. B is P.D. Then, we solve (2) in an approximate sense in a 2D subspace:

$$p_* = \arg\min m_k(p)$$

subject to:

- $\|p\|_2 \leq \Delta$
- $\underline{p} \in \text{Span}(\underline{g}, B^{-1}\underline{g})$.

i.e.

$$p_* = \min_{\alpha, \beta} m_k(\overbrace{\alpha \underline{g} + \beta^{-1} \underline{g}}^{p})$$

subject to $\|p\|_2^2 \leq \Delta^2$.

α, p

subject to $\|p\|_2^2 \le \Delta^2$.

2. When B has a zero eigenvalue but no negative eigenvalue, take

$$p_* = p_{\text{Cauchy}}$$

3. When B has negative eigenvalues, we take

$$p_* = \arg\min \; m_k(p)$$

subject to :

- $\|p\|_2 \le \Delta$
- $p \in \operatorname{span}\left(g, (B + \alpha \mathbb{I})^{-1} g\right)$

such that $B + \alpha\mathbb{I}$ has only positive eigenvalues, thus

$$\alpha \in (-\lambda_1, -2\lambda_1]$$

where λ_1 is the most negative eigenvalue of B. i.e. a perturbation of B.

Chapter 18 – Global Optimization via Simulated Annealing

Motivation : Methods based on a descent direction will get stuck at a local minimum

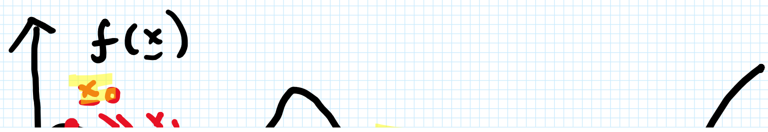

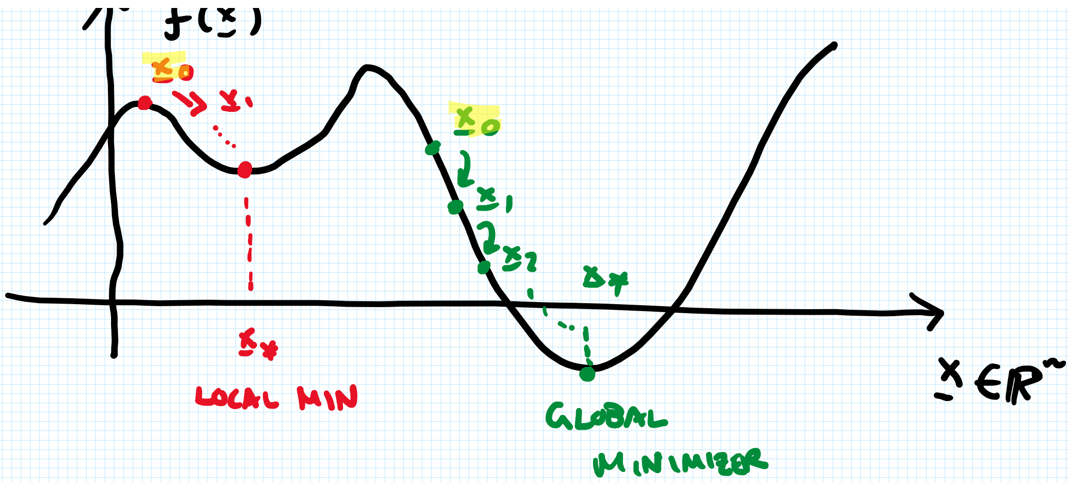

We want to look at methods that converge to the global minimizer, regardless of the starting-point.

Global Optimization :

- Simulated Annealing
- Genetic Algorithms
- Particle Swarm Optimization

} Metaheuristic algorithms

We look at the theory behind Simulated Annealing.

Advantages:

- Converges to global min
- No differentiation of cost function required, hence cost function does not have to be smooth

Disadvantages:

- Slow compared to methods based on a descent direction.

Simulated Annealing (S.A.) is inspired by Physics. Suppose we have a collection of particles (=system) with n degrees of freedom. The energy of the system is a map

$$E : \mathbb{R}^n \longrightarrow \mathbb{R} \qquad (1)$$
$$\underline{x} \longmapsto E(\underline{x})$$

We want to know the probability $d\mathbb{P}$ of finding the system in a small volume $d^n x$ of phase space, centred at $\underline{x}$:

$$d\mathbb{P} = p(\underline{x})\, d^n x \qquad (2)$$

$p(\underline{x})$ is the probability distribution :

$$\left.\begin{array}{l} \bullet \quad p(\underline{x}) \geq 0 \\ \bullet \quad \int_{\mathbb{R}^n} p(\underline{x}) d^n x = 1 \end{array}\right\} \qquad (3)$$

Hence, the energy of the system is :

$$\bar{E} = \int_{\mathbb{R}^n} E(\underline{x})\, p(\underline{x})\, d^n x \qquad (4)$$

Entropy of the system :

$$S = - \int_{\mathbb{R}^n} p(\underline{x}) \log p(\underline{x})\, d^n x \qquad (5)$$

What is $p(\underline{x})$? Solution: That distribution which maximizes S, for a fixed energy ("Canonical Ensemble")

$$\tilde{S} = -\int p(\underline{x}) \log p(\underline{x})\, d^n x$$

$$- \beta \left(\int E(\underline{x}) \log p(\underline{x})\, d^n x - \bar{E} \right)$$

$$- \alpha \left(\int p(\underline{x})\, d^n x - 1 \right)$$

(integrals over all of $\mathbb{R}^n$).

We look at small variations in $\tilde{S}$ obtained by taking small variations in p ($p + \delta p$).

$$\delta \tilde{S} = - \int \delta (p \log p)\, d^n x$$

$$- \beta \left(\int E(\underline{x})\, \delta p\, d^n x - \cancel{E} \right)$$

$$\underline{+ \alpha} \left(\int \delta p\, d^n x \right)$$

$$\delta(p \log p) = \delta p \cdot \log p + p \cdot \delta \log p$$

$$= \delta p \cdot \log p + \cancel{p} \frac{\delta p}{\cancel{p}}$$

$$= \delta p (\log p + 1)$$

Hence:

$$\delta \tilde{S} = - \int \delta p (\log p + 1)\, d^n x$$

$$- \beta \int E(\underline{x}) \, \delta p \, d^n x$$

$$+ \quad \alpha \int \delta p$$

$$= \quad - \int \delta p \left[\log p + 1 + \beta E(\underline{x}) - \alpha \right] d^n x$$

When p is at the stationary value, $\delta \tilde{S} = 0$:

$$\int \delta p \left[\dots \right] d^n x = 0 . \quad \forall \, \delta p$$

Hence, $[\dots] = 0$.

$$\Rightarrow \log p + \beta E + 1 - \alpha = 0$$

$$\Rightarrow \quad p(\underline{x}) = e^{-\beta E(\underline{x})} \; e^{\alpha - 1}$$

$$\Rightarrow \quad p(\underline{x}) = \frac{e^{-\beta E(\underline{x})}}{\int e^{-\beta E(\underline{x})} \, d^n x}$$

$$Z = \int e^{-\beta E(\underline{x})} \, d^n x$$, PARTITION FUNCTION

$$p(\underline{x}) = \frac{e^{-\beta E(\underline{x})}}{Z}$$ BOLTZMAN DISTRIBUTION

Sub back into $\tilde{S}$:

$$\tilde{S}\left[p(\underline{x}) = \frac{e^{-\beta E(\underline{x})}}{Z}\right] = \beta \bar{E} + \log Z$$

$$\Rightarrow \boxed{S_{max} = \beta \bar{E} + \log Z}$$

$$\beta = \frac{\partial S_{max}}{\partial \bar{E}}$$

$$\Rightarrow \boxed{\frac{1}{T} = \frac{\partial S_{max}}{\partial \bar{E}}}$$

$$\frac{\partial}{\partial \beta} \log Z = \frac{1}{Z}\frac{\partial Z}{\partial \beta}$$

$$= \frac{1}{Z}\frac{\partial}{\partial \beta}\int e^{-\beta E(\underline{x})}\, d^n x$$

$$= \frac{1}{Z}\int \frac{\partial}{\partial \beta} e^{-\beta E(\underline{x})}\, d^n x$$

$$= \frac{1}{Z}\int -E(\underline{x})\, e^{-\beta E(\underline{x})}\, d^n x$$

$$= -\int E(\underline{x})\, \underbrace{\frac{e^{-\beta E(\underline{x})}}{Z}}\, d^n x$$

$$= -\bar{E}$$

$$\boxed{\bar{E} = -\frac{\partial \log Z}{\partial \beta}}$$

<u>The Quench</u> : Suppose that $E(\underline{x})$ has a global minimizer at $\underline{x}_*$, with

a global minimizer at $\underline{x}_*$, with

$$E(\underline{x}_*) = E_{min}.$$

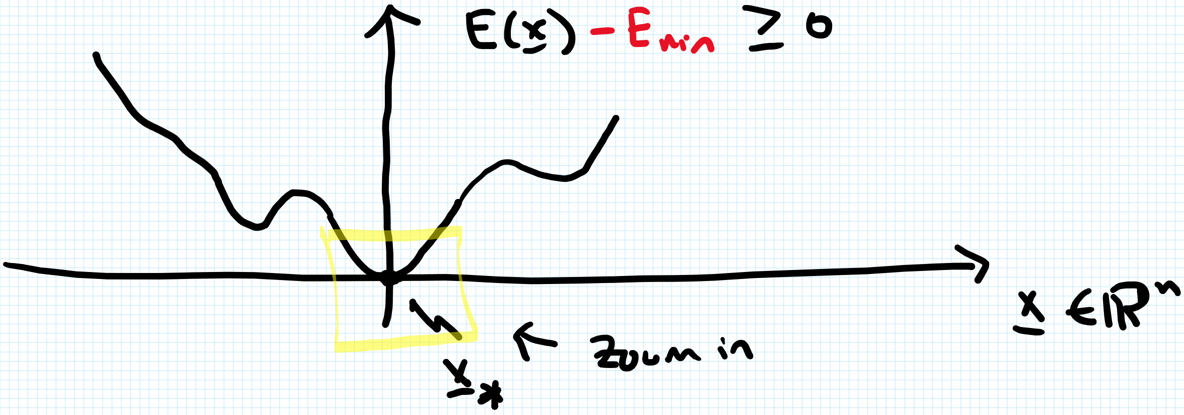

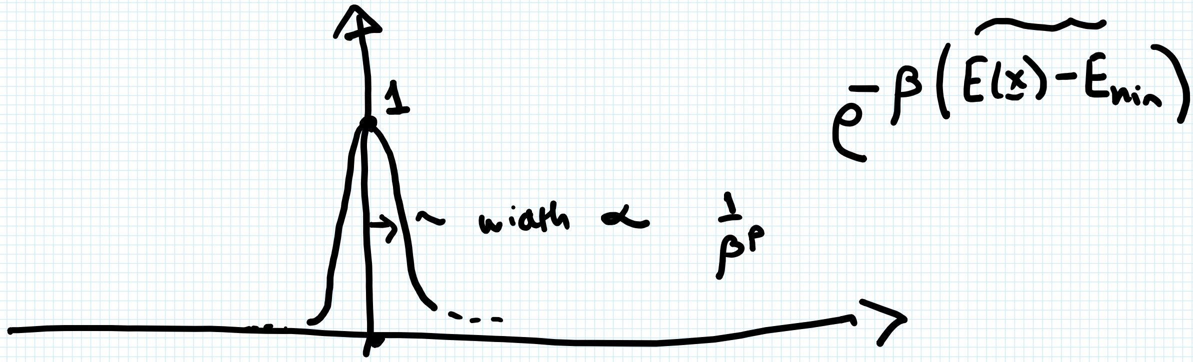

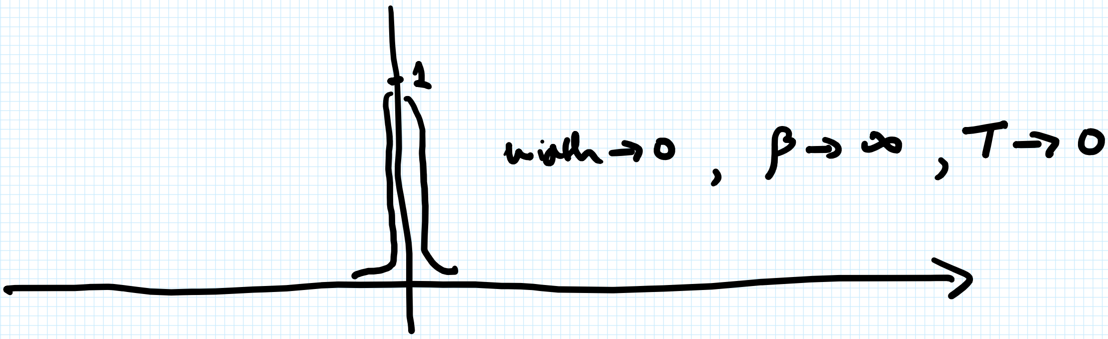

$$\frac{e^{-\beta(E(\underline{x}) - E_{min})}}{Z} \longrightarrow \delta(\underline{x} - \underline{x}_*)$$

As the system is cooled ("quenched"), $p(\underline{x}) \longrightarrow \delta(\underline{x} - \underline{x}_*)$, so the probability of being in the minimum-energy state tends to one.

of being in the minimum-energy state tends to one.

So, if we can simulate a collection of states of the system,

$$\underline{x}_0, \underline{x}_1, \ldots, \underline{x}_n$$

and at the same time reduce the temperature, then in the limit

$$\underline{x}_n \to \underline{x}_* \quad \text{as} \quad n \to \infty$$

Hence, $\underline{x}_* = \lim_{n\to\infty} \underline{x}_n$ & (global minimizer)

can be computed from an iterative process.

Plan for next week:

- Tuesday: Online lecture / F2F, convergence proof of SA algorithm (EXAM)
- Thursday:
  - Coding exercise - Monte Carlo simulation EXAM PREP
  - Structure of Exam # 1

$6N = n.$

$\mathbb{R}^n$