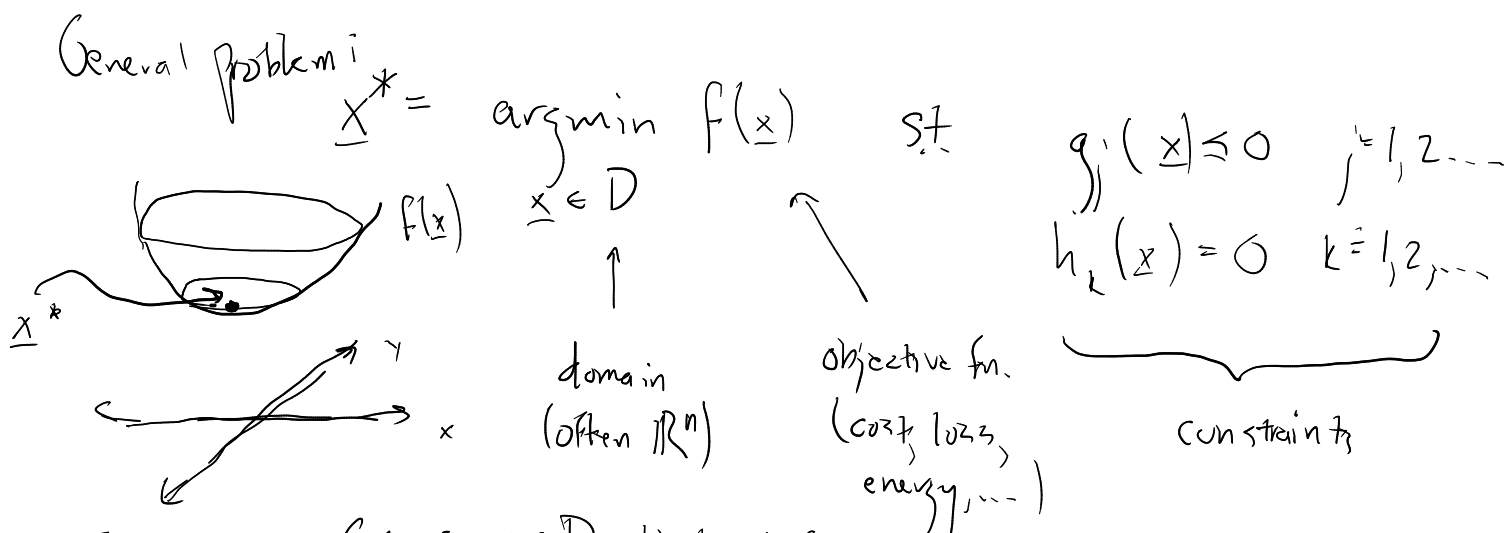


Optimization

- Topics:
- 1) Problem definition, types of problems
 - 2) Convex problems
 - 3) Solution methods
 - 4) SVMs

Boyd & Vandenberghe
Convex Optimization

Green notes are extra examples and extensions not covered in class.



Feasible set: Set of $\underline{x} \in D$ that satisfy constraints.

Brute force (grid search): Divide D into grid w/ width δ . # points grows as $(1/\delta)^n$

Type	Domain	Objective	Constraints	Solution
Linear	\mathbb{R}^n	$\underline{c}^T \underline{x}$	$A \underline{x} \leq \underline{b}, \underline{x} \geq 0$	Easy (simple method)
Integer	\mathbb{N}^n	" "	" "	NP hard in general
Constraint set	$\{0, 1\}^n$	Constant	Boolean	NP hard in general
Convex	\mathbb{R}^n	Convex fn.	Convex set	Easy! (interior point method)
Quadratic	\mathbb{R}^n	$\underline{x}^T Q \underline{x} + \underline{c}^T \underline{x}$	$A \underline{x} \leq \underline{b}$	Easy if Q positive-definite
...				

Ex (least squares) $y = X \cdot \beta$. Given $\{x_i, y_i\}, i=1, \dots, P$, find optimal β

$$\beta^* = \underset{\beta}{\operatorname{argmin}} \| X\beta - y \|^2$$

$\uparrow \quad \uparrow \quad \uparrow$
 $P \times N \quad N \times 1 \quad P \times 1$

$$f(\beta) = (X\beta - y)^T (X\beta - y)$$

$$= \beta^T X^T X \beta - 2y^T X \beta + \underbrace{y^T y}_{\text{constant, ignore}}$$

$$\beta^* = \underset{\beta}{\operatorname{argmin}} \frac{1}{2} \beta^T Q \beta + c^T \beta, \quad Q = X^T X$$
$$c = -2y^T X$$

Quadratic problem (convex)

Regularization:

$$f(\beta) = \underbrace{(X\beta - y)^T (X\beta - y)}_{\text{fitting to data}} + \underbrace{\lambda \beta^T \beta}_{\text{penalty on large } \beta_i^2}$$


Same as above with

$$Q \leftarrow Q + \lambda I$$


$$= \lambda \beta^T I \beta$$

\underline{x}^* is global optimum. May be local optima \underline{x}^l s.t.

$$f(\underline{x}^l) \leq f(\underline{x}) \text{ for } \|\underline{x} - \underline{x}^l\| < \varepsilon.$$




Def: $f(x)$ is convex if $f(ta + (1-t)b) \leq tf(a) + (1-t)f(b)$



\Rightarrow If $\nabla f(a)$ is slope at a ,

$$f(a) + \nabla f(a)[x - a] < f(x)$$



If f convex, all local min. are global min.



strictly convex, one global min.

Ex: $f(x) = x^2$. $f(ta + (1-t)b) = t^2 a^2 + (1-t)^2 b^2 + 2t(1-t)ab$ (1)

$$tf(a) + (1-t)f(b) = ta^2 + (1-t)b^2$$
 (2)

$$(1)-(2) = (t^2 - t)a^2 + ((1-t)^2 - (1-t))b^2 + 2t(1-t)ab$$

$$= t(t-1)a^2 + t(t-1)b^2 - 2t(t-1)ab = t(t-1)(a-b)^2 \leq 0$$



Higher d:

Gradient of $f(x)$: $\nabla f(x) = \begin{pmatrix} \partial f / \partial x_1 \\ \partial f / \partial x_2 \\ \vdots \\ \partial f / \partial x_n \end{pmatrix}$

Hessian:

$$H(x) = \begin{pmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & & \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2^2} & & \\ & & \ddots & \\ & & & \ddots \end{pmatrix}$$

1d: $f(x) \approx f(a) + f'(a)(x-a) + \frac{1}{2}f''(a)(x-a)^2$

Minimum: $f' = 0, f'' > 0.$

Strictly convex if $f'' > 0$ everywhere.

Higher-d: $f(x) \approx f(a) + \nabla f(a)(x-a) + \frac{1}{2}(x-a)^T H(a)(x-a)$

Minimum: $\nabla f = 0, H$ positive definite ($v^T H v > 0 \forall v$)

Strictly convex if H positive definite everywhere

Analytical approaches:

Unconstrained problem: look for \underline{x}^* w/ $\nabla f(\underline{x}^*) = 0$

Equality constraints \rightarrow method of Lagrange multipliers.

$$\min_{\underline{x}} f(\underline{x}) \quad \text{s.t.} \quad g_j(\underline{x}) = 0 \quad j = 1, 2, \dots$$



$$\text{let } \mathcal{L}(\underline{x}, \underline{\lambda}) = f(\underline{x}) - \sum_i \lambda_i g_i(\underline{x})$$

λ_i are Lagrange multipliers.

Look for \underline{x}^* , $\underline{\lambda}$ s.t. $\nabla \mathcal{L}(\underline{x}, \underline{\lambda}) = 0$

$$\text{Note } \frac{\partial}{\partial \lambda_i} \mathcal{L} = -g_i(\underline{x}) = 0$$

$$\nabla_{\underline{x}} f(\underline{x}) - \sum_i \lambda_i \nabla_{\underline{x}} g_i(\underline{x}) = 0 \quad \text{"level sets"}$$




With inequality constraints $g_j(x) \leq 0$, $h_k(x) = 0$.
 Karush-Kuhn Tucker (KKT) conditions

Numerical approaches:

Linear problems:

$$\min \underline{c}^T \underline{x} \quad \text{s.t.} \quad A \underline{x} \leq \underline{b}, \quad \underline{x} \geq 0$$



$\downarrow \underline{c}$

Simplex method.


Convex problems: Interior point methods.

Gradient descent:

Given initial value \underline{x}_0 ,

$$\underline{x}_{n+1} = \underline{x}_n - \eta_n \nabla f(\underline{x}_n)$$


\uparrow
 step size
 "learning rate"



Often $\eta_{n+1} < \eta_n$.

"line search": Choose η_n to minimize $f(\underline{x}_{n+1})$.

Problems: 1) multiple minima (multiple initial conditions, noise)
2) small gradients



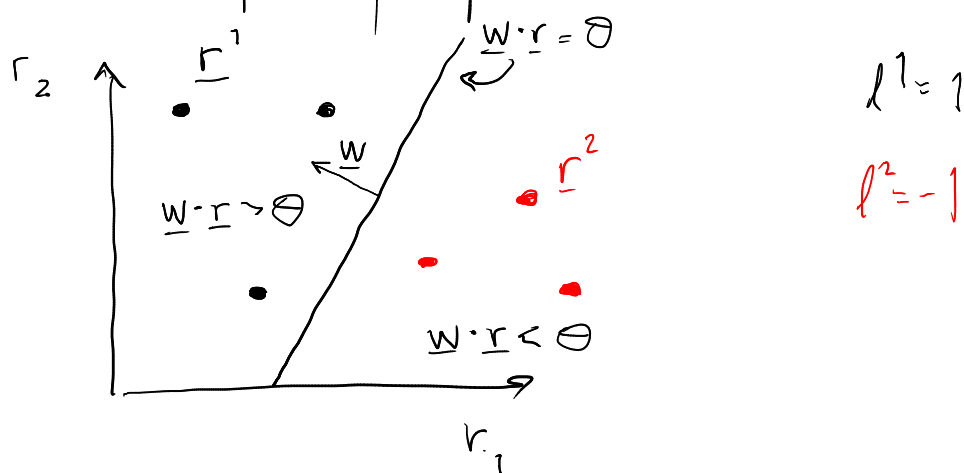
Momentum: $\underline{z}_{n+1} = \beta \underline{z}_n + \nabla f(\underline{x}_n)$

$$\beta = 0.99$$

$$\underline{x}_{n+1} = \underline{x}_n - \eta_n \underline{z}_{n+1}$$

Convex optimization example: Support vector machines (SVM)

Problem: Given P patterns \underline{r}^u , $u=1, \dots, P$, and labels $l^u = \pm 1$, find linear separating hyperplane that optimally separates $l=1$ and $l=-1$.




Classifier:
$$l = \text{sign} \left(\underbrace{\underline{w} \cdot \underline{r}}_{\text{weights}} - \underbrace{\theta}_{\text{threshold}} \right)$$

How to choose optimal \underline{w} , θ ? May be multiple valid solutions (draw).

Idea: Maximize margin m^u (smallest distance from \underline{r}^u to boundary). (draw).

If $\|\underline{w}\|=1$, then
$$m^u = \left| \hat{\underline{w}} \cdot \underline{r}^u - \theta \right|$$



\underline{w} is perpendicular to separating hyperplane.

Optimization Problem:

$$\underset{\underline{w}}{\text{maximize}} \quad \min_u \left| \underline{w} \cdot \underline{r}^u - \theta \right| \quad \text{s.t.} \quad \|\underline{w}\| = 1, \\ \text{sign}(\underline{w} \cdot \underline{r}^u - \theta) = l^u$$

Redefine constraints: $(\underline{w} \cdot \underline{r}^u - \theta) \cdot l^u > 0$.

How to deal with $\|\underline{w}\|$?

If margin = m , $(\underline{w} \cdot \underline{r}^u - \theta) l^u \geq m$. Divide by m :

$$\left(\frac{\underline{w}}{m} \cdot \underline{r}^u - \frac{\theta}{m} \right) l^u \geq 1$$

Reparameterize: $\tilde{\underline{w}} \leftarrow \frac{\underline{w}}{m}$
 $\tilde{\theta} = \frac{\theta}{m}$

$$(\tilde{\underline{w}} \cdot \underline{r}^u - \tilde{\theta}) l^u \geq 1.$$

Note: $\|\underline{\tilde{w}}\| = \left\| \frac{\underline{w}}{m} \right\| = \frac{1}{m}$

Maximize margin \iff minimize $\|\underline{\tilde{w}}\|^2!$

$$\underline{w}^*, \theta^* = \underset{\underline{w}, \theta}{\operatorname{argmin}} \quad \underline{w}^T \underline{w} \quad \text{s.t.} \quad (\underline{w} \cdot \underline{r}^u - \theta) l^u \geq 1.$$

$$\implies \|\underline{w}^*\| = \frac{1}{m}$$

Properties: 1) \underline{w} determined only by closest points (those on margin) — support vectors.

2) Binary classification — linear boundary

3) Fully supervised (l^u known $\forall u$)

4) Sol'n only exists if data linearly separable

Extensions: 1) Multi-class

2) Kernel SVM (later)


3) "Soft margin"

↓
Allow misclassifications:

Penalize w/ $c^u = \max(0, 1 - l^u(\underline{w} \cdot \underline{r}^u - \theta))$

$$\min \sum_{u=1}^P c^u + \lambda \|\underline{w}\|^2$$

"hinge loss"



How to write as convex problem?

Rewrite constraints: $(\underline{w} \cdot \underline{r}^m - \Theta) l^m \geq 1 - c^m$

$$\underline{w}^*, \underline{c}^*, \Theta^* = \underset{\underline{w}, \underline{c}, \Theta}{\operatorname{argmin}} \sum_{m=1}^p c^m + \lambda \underline{w}^T \underline{w} \quad \text{s.t.} \quad \begin{aligned} (\underline{w} \cdot \underline{r}^m - \Theta) l^m \\ \geq 1 - c^m \end{aligned}$$

Interpretation of Lagrange mult:

$$\mathcal{L} = f(\underline{x}) + \lambda g(\underline{x}). \quad \text{let } g(\underline{x}) = \hat{g}(\underline{x}) - c = 0$$

$\Rightarrow \frac{\partial \mathcal{L}}{\partial c} = \lambda$. λ is sensitivity of \mathcal{L} to change in constant

For SVMs, $\lambda \neq 0$ only for SVs (on margin).

Duality: "primal problem"

$$\underline{x}^* = \underset{\underline{x}}{\operatorname{argmin}} f(\underline{x}), \quad \begin{aligned} g_j(\underline{x}) &\leq 0 \\ h_k(\underline{x}) &= 0 \end{aligned}$$

Write Lagrangian

$$\mathcal{L}(\underline{x}, \underline{\lambda}, \underline{\nu}) = f(\underline{x}) + \sum_j \lambda_j g_j(\underline{x}) + \sum_k \nu_k h_k(\underline{x})$$

Dual function: $\approx \min$

$$G(\underline{\lambda}, \underline{\nu}) = \inf_{\underline{x} \in D} \mathcal{L}(\underline{x}, \underline{\lambda}, \underline{\nu})$$

Note $G(\underline{\lambda}, \underline{v}) \leq f^*$ if $\lambda_j \geq 0 \forall j$.

Why? For any feasible \underline{x} , $f_k(\underline{x}) \leq 0$, $h_k(\underline{x}) = 0$

$$\Rightarrow \underbrace{\sum_j \lambda_j f_j(\underline{x})}_{\leq 0} + \underbrace{\sum_k v_k h_k(\underline{x})}_{= 0} \leq 0$$

$$\Rightarrow \mathcal{L}(\underline{x}, \underline{\lambda}, \underline{v}) \leq f(\underline{x})$$

Can minimize G "dual problem":

$$\underline{\lambda}^*, \underline{v}^* = \underset{\underline{\lambda}, \underline{v}}{\text{argmax}} G(\underline{\lambda}, \underline{v}) \quad \text{s.t. } \lambda_j \geq 0$$

If primal problem convex, $G^* = f^*$

For SVM,

$$\mathcal{L} = \frac{1}{2} \underline{w}^T \underline{w} - \sum_{\mu} \lambda^{\mu} \left[(\underline{w} \cdot \underline{r}^{\mu} - \theta) l^{\mu} - 1 \right]$$

$$G(\underline{\lambda}) = \underset{\underline{w}, \theta}{\text{inf}} \mathcal{L} \quad \frac{\partial \mathcal{L}}{\partial w_i} = 0 \Rightarrow \underline{w} - \sum_{\mu} \lambda_{\mu} \underline{r}^{\mu} l^{\mu} = 0$$

$$\frac{\partial \mathcal{L}}{\partial \theta} = 0 \Rightarrow \sum_{\mu} \lambda^{\mu} l^{\mu} = 0$$

$$\underline{w} = \sum_{\mu} \lambda_{\mu} \underline{r}^{\mu} l^{\mu} \leftrightarrow \text{sum of SVs}$$

Dual problem:

$$\max \frac{1}{2} \sum_u \sum_v \lambda_u \lambda_v d_u d_v (\underline{r}^u)^T \underline{r}^v + \sum_u \lambda^u$$

$$\text{s.t. } \sum_u \lambda_u d_u = 0, \quad \lambda_u \geq 0.$$

Optimization over P variables vs. n .