Convex Optimization

Chapter 2 – Convex Sets

- *Basics*
  - A set is *affine* if it contains any line through two of its point. Alternatively, $\boldsymbol{x}_1,\cdots,\boldsymbol{x}_n \in \mathcal{C}, \theta_+ = 1 \Rightarrow \theta_1\boldsymbol{x}_1 + \cdots + \theta_n\boldsymbol{x}_n \in \mathcal{C}$.
  - The *affine hull* of a set of points is the set of all affine combinations of these points.
  - The *affine dimension* of a set is the dimension of its affine hull. Its *relative interior* is its interior relative to its affine hull

$$\text{relint } \mathcal{C} = \left\{\boldsymbol{x} \in C : B(\boldsymbol{x}, r) \cap \text{aff } \mathcal{C} \subseteq \mathcal{C} \text{ for some } r > 0\right\}$$

  - The most general form of a convex combination is $\mathbb{E}(\boldsymbol{x})$, where $\mathbb{P}(\boldsymbol{x} \in \mathcal{C}) = 1$.
  - A set $\mathcal{C}$ is a *cone* if $\boldsymbol{x} \in \mathcal{C}, \theta \geq 0 \Rightarrow \theta\boldsymbol{x} \in \mathcal{C}$
    - The set $\left\{(\boldsymbol{x}, t) : \left\|\boldsymbol{x}\right\| \leq t\right\}$ is a *norm cone* associated with a particular norm.
    - The conic hull of $\{\boldsymbol{x}_i\}$ is $\left\{\lambda_1\boldsymbol{x}_1 + \cdots + \lambda_k\boldsymbol{x}_k : \boldsymbol{\lambda} \geq \boldsymbol{0}\right\}$.
  - A hyperplane is a set of the form $\left\{\boldsymbol{x} : \boldsymbol{a} \cdot \boldsymbol{x} = b\right\}$. Hyperplane with normal vector $\boldsymbol{a}$, offset $\boldsymbol{b}$ from the origin; can be written as $\left\{\boldsymbol{x} : \boldsymbol{a} \cdot (\boldsymbol{x} - \boldsymbol{x}_0) = 0\right\} = \boldsymbol{x}_0 + \boldsymbol{a}^\perp$
  - Given $k + 1$ *affinely independent pints* (ie: $\boldsymbol{v}_i - \boldsymbol{v}_0$ linearly independent), the *k-dimensional simplex* determined by these points is $\mathcal{C} = \left\{\sum \theta_i \boldsymbol{v}_i : \theta_i \geq 0, \theta_+ = 1\right\}$. We can describe this as a polyhedron as follows:
    - Write $B = \left[\boldsymbol{v}_1 - \boldsymbol{v}_0, \cdots, \boldsymbol{v}_k - \boldsymbol{v}_0\right]$ and $\boldsymbol{\theta}' = \left[\theta_1, \cdots, \theta_k\right]$. All points $x \in \mathcal{C}$ can then be expressed as $\boxed{\boldsymbol{x} = \boldsymbol{v}_0 + B\boldsymbol{\theta}'}$ provided $\boldsymbol{\theta}' \geq 0$ and $\boldsymbol{1} \cdot \boldsymbol{\theta}' \leq 1$
    - B has rank k (by assumptions) and $k \leq n$, and so there exists a $A \in \mathbb{R}^{n \times n}$ such that $AB = \begin{bmatrix} A_1 B \\ A_2 B \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix}$.

- Multiplying the boxed equation by A, we get $\boldsymbol{\theta}' = A_1\boldsymbol{x} - A_1\boldsymbol{v}_0$ and $A_2\boldsymbol{x} = A_2\boldsymbol{v}_0$. We can therefore express $\boldsymbol{\theta}' \geq 0$ and $\boldsymbol{1}^\top\boldsymbol{\theta}' \leq 0$ as linear inequalities. Together with $A_2\boldsymbol{x} = A_2\boldsymbol{v}_0$, they define the polyhedron.

- ***Operations that preserve convexity***
  - ***Intersection*** (including infinite intersection) – also preserve subspaces, affine sets and convex cones:
    - ***Example***: The positive semidefinite cone $\mathbb{S}_n^+$ can be written as $\bigcap_{z\neq 0}\left\{X \in \mathbb{S}_n : \boldsymbol{z}^\top X\boldsymbol{z} \geq 0\right\}$. Each set in the intersection is convex (since the defining equations are linear), and so $\mathbb{S}_n^+$ is convex. □
    - ***Example***: $\mathcal{S} = \left\{\boldsymbol{x} \in \mathbb{R}^m : \left|\sum x_i \cos(it)\right| \leq 1 \text{ for } t \in [-\frac{\pi}{3}, \frac{\pi}{3}]\right\}$ can be written as $\bigcap_{t\in[-\frac{\pi}{3},\frac{\pi}{3}]}\left\{X \in \mathbb{S}_n : -1 \leq (\cos t, \cdots, \cos mt)\cdot \boldsymbol{x} \leq 1\right\}$, and so is convex. □
  - ***Affine functions***: An affine function has the form $f(\boldsymbol{x}) = A\boldsymbol{x} + \boldsymbol{b}$. The image and inverse image of a convex set under such a function is convex.
    - ***Example***: $\mathcal{S}_1 + \mathcal{S}_2 = \left\{\boldsymbol{x} + \boldsymbol{y} : \boldsymbol{x} \in \mathcal{S}_1, \boldsymbol{y} \in \mathcal{S}_2\right\}$ is the image of $\mathcal{S}_1 \times \mathcal{S}_2 = \left\{(\boldsymbol{x}_1, \boldsymbol{x}_2) : \boldsymbol{x}_1 \in \mathcal{S}_1, \boldsymbol{x}_2 \in \mathcal{S}_2\right\}$ under $f(\boldsymbol{x}_1, \boldsymbol{x}_2) = \boldsymbol{x}_1 + \boldsymbol{x}_2$. □
    - ***Example***: $\left\{\boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}, C\boldsymbol{x} = \boldsymbol{d}\right\}$ is the inverse image of $\mathbb{R}_+ \times \{0\}$ under $f(\boldsymbol{x}) = (\boldsymbol{b} - A\boldsymbol{x}, \boldsymbol{d} - C\boldsymbol{x})$. □
    - ***Example***: $\left\{\boldsymbol{x} : A(\boldsymbol{x}) = \boldsymbol{x}_1 A_1 + \cdots + \boldsymbol{x}_n A_n \preceq B\right\}$ is the inverse image of the positive semidefinite cone $\mathbb{S}_+^n$ under $f(\boldsymbol{x}) = B - A(\boldsymbol{x})$. □
    - ***Example***: $\left\{\boldsymbol{x} : (\boldsymbol{x} - \boldsymbol{x}_c)^\top P(\boldsymbol{x} - \boldsymbol{x}_c) \leq 1\right\}$, where $P \in \mathbb{S}_{++}^n$ is the image of a unit Euclidean ball under $f(\boldsymbol{u}) = P^{1/2}\boldsymbol{u} + \boldsymbol{x}_c$. □
  - ***Perspective function***: $f(\boldsymbol{z}, t) = \boldsymbol{z} / t$, where $t > 0$. It normalizes the last component of a vector to 1 and then gets rid of that component. The image of a convex set under the perspective function is convex.
  - ***Linear-fractional function***: A linear-fractional function is formed by compsing that perspective function with an affine function. They take the form $f(\boldsymbol{x}) = (A\boldsymbol{x} + b) / (\boldsymbol{c}\cdot\boldsymbol{x} + d)$, with domain $\left\{\boldsymbol{x} : \boldsymbol{c}\cdot\boldsymbol{x} + d > 0\right\}$.

- ***Separating & Supporting Hyperplanes***
  - ***Theorem***: If $\mathcal{C} \cap \mathcal{D} \neq \varnothing$ then $\exists \boldsymbol{a} \neq \boldsymbol{0}$ and b such that $\boldsymbol{a} \cdot \boldsymbol{x} \leq \boldsymbol{b}\ \forall \boldsymbol{x} \in \mathcal{C}$ and $\boldsymbol{a} \cdot \boldsymbol{x} \geq \boldsymbol{b}\ \forall \boldsymbol{x} \in \mathcal{D}$. In some cases, *strict* separation is possible (ie: the inequalities become strict).
  - ***Example***: Consider an affine set $\mathcal{D} = \left\{ F\boldsymbol{u} + \boldsymbol{g} : \boldsymbol{u} \in \mathbb{R}^m \right\}$ and a convex set $\mathcal{C}$ which are disjoint. Then by our Theorem, there exists $\boldsymbol{a} \neq \boldsymbol{0}$ and $\boldsymbol{b}$ such that $\boldsymbol{a} \cdot \boldsymbol{x} \leq \boldsymbol{b}\ \forall \boldsymbol{x} \in \mathcal{C}$ and $\boldsymbol{a} \cdot [F\boldsymbol{u} + \boldsymbol{g}] \geq \boldsymbol{b} \Rightarrow \boldsymbol{a}^\top F\boldsymbol{u} \geq \boldsymbol{b} - \boldsymbol{a} \cdot \boldsymbol{g}\ \forall \boldsymbol{u}$. The only way a linear function can be bounded below is if it's 0 – as such, $\boldsymbol{a}^\top F = \boldsymbol{0}$, and $\boldsymbol{b} \leq \boldsymbol{a} \cdot \boldsymbol{g}$.
  - ***Theorem***: Consider two convex sets $\mathcal{C}$ and $\mathcal{D}$. Provided at least one of them is open, they are disjoint if and only if there exists a separating hyperplane.

    Proof: Consider the open set – $\boldsymbol{a} \cdot \boldsymbol{x}$ cannot be 0 for any $\boldsymbol{x}$ in that set, else it would be greater than 0 for a point close to $\boldsymbol{x}$. Thus, $\boldsymbol{a} \cdot \boldsymbol{x}$ is *strictly* less than 0 for *all* points in the open set. ■
  - ***Example***: Consider $A\boldsymbol{x} < \boldsymbol{b}$. This has a solution if and only if $\mathcal{C} = \left\{ \boldsymbol{b} - A\boldsymbol{x} : \boldsymbol{x} \in \mathbb{R}^n \right\}$ and $\mathcal{D} = \mathbb{R}^m_{++}$ do not intersect. By the Theorem, this is true if and only if there exists $\boldsymbol{\lambda} \neq \boldsymbol{0}$ and $\mu \in \mathbb{R}$ such that $\boldsymbol{\lambda} \cdot \boldsymbol{y} \leq \mu\ \forall \boldsymbol{y} \in \mathcal{C}$ and $\boldsymbol{\lambda} \cdot \boldsymbol{y} \geq \mu\ \forall \boldsymbol{y} \in \mathcal{D}$. In other words, there is *not* separating hyperplane iff

    $$\boldsymbol{\lambda} \neq \boldsymbol{0} \qquad \boldsymbol{\lambda} \geq \boldsymbol{0} \qquad A^\top \boldsymbol{\lambda} = \boldsymbol{0} \qquad \boldsymbol{\lambda}^\top \boldsymbol{b} \leq 0$$

    Thus, only one of this sytem and $A\boldsymbol{x} < \boldsymbol{b}$ can have a solution. □

Chapter 3 – Convex Functions

- ***Basics***
  - We extend a convex/concave function by setting it to $+/-\infty$ outside its domain.
  - ***Theorem***: f is convex over $\mathcal{C}$ convex iff $f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x})$ over $\mathcal{C}$.

    Proof: $\boxed{\Leftarrow}$ choose x_1, x_2 and convex comb z. Apply equation with $y = z$ and $x = x_i$. Multiply one equation by λ, other by $1 - \lambda$. Add the two. $\boxed{\Rightarrow}$ Take x, y. By convexity $f(x + t[y - x]) \leq (1 - t)f(x) + tf(y)$ for $t \in (0,1)$. Re-arrange to get

$f(y)$ on one side, divide by t, take limit as $t \to 0$. General case consider $g(t) = f\left(t\boldsymbol{y} + (1-t)\boldsymbol{x}\right)$ and $g'(t) = \nabla f\left(t\boldsymbol{y} + (1-t)\boldsymbol{x}\right)^\top \left(\boldsymbol{y} - \boldsymbol{x}\right)$. $\Rightarrow$ Apply previous result with $y = 1$ and $x = 0$. $\Leftarrow$ Apply inequality with $t\boldsymbol{y} + (1-t)\boldsymbol{x}$ and $\tilde{t}\boldsymbol{y} + (1-\tilde{t})\boldsymbol{x}$. This implies an inequality about g that makes it convex. ■

- ***Theorem***: $\nabla^2 f(\boldsymbol{x}) \succeq 0$ over $\mathcal{C}$ convex $\Rightarrow f$ convex over $\mathcal{C}$.

  Proof: $f(\boldsymbol{y}) = f(\boldsymbol{x}) + \nabla f(\boldsymbol{x}) \cdot (\boldsymbol{y} - \boldsymbol{x}) + \frac{1}{2}(\boldsymbol{y} - \boldsymbol{x})^\top \left[\nabla^2 f\left(\varepsilon \boldsymbol{x} + (1-\varepsilon)\boldsymbol{y}\right)\right](\boldsymbol{y} - \boldsymbol{x})^\top$ for $\varepsilon \in [0,1]$. If $\nabla^2 f$ is positive definite, get FOC for convexity. ■

Convex functions

The following functions are convex

Function	**Parameters**	**Convex/concave…**	**…on domain**
e^{ax}	$a \in \mathbb{R}$	convex	$\mathbb{R}$
x^a	$a \geq 1$ or $a \leq 0$	convex	$(0, \infty)$
	$0 \leq a \leq 1$	concave	$(0, \infty)$
$\lvert x \rvert^p$	$p \geq 1$	convex	$\mathbb{R}$
$\log x$		concave	$(0, \infty)$
$x \log x$		convex	$(0, \infty)$
$\boldsymbol{a} \cdot \boldsymbol{x} + \boldsymbol{b}$	(ie: any affine function)	both	$\mathbb{R}^n$
$\lVert \; \rVert$	(ie: any norm)	convex	$\mathbb{R}^n$
$\log\left(\sum e^{x_i}\right)$	(the log-sum-exp func.)	convex	$\mathbb{R}^n$
$\left(\prod x_i\right)^{1/n}$	(the geometric mean)	concave	$(0, \infty)^n$
$\log \det X$	(the log determinant)	convex	$X \in \mathbb{S}^n_{++}$
$\sum w_i f_i(\boldsymbol{x})$	$w_i \geq 0$	Same as f_i, providing they are all concave/convex	

Ways to find convexity:

- ***Directly verify definition***
- ***Check the Hessian***: for example, for $f(x,y) = x^2 / y$

$$\nabla^2 f(x,y) = \frac{2}{y^3}\begin{bmatrix} y^2 & -xy \\ -xy & x^2 \end{bmatrix} = \frac{2}{y^3}\begin{bmatrix} y & -x \end{bmatrix}\begin{bmatrix} y \\ -x \end{bmatrix} \succeq 0$$

- ***Restrict to a line***: f is convex if and only if $g(t) = f\left[t\boldsymbol{x}_1 + (1-t)\boldsymbol{x}_2\right]$ is convex over $[0,1]$ $\forall\ \boldsymbol{x}_1, \boldsymbol{x}_2$. For example, $f(X) = \log\det X$. Take the line $X = Z + tV$, restricting to values of t for which $X \succ 0$, and wlog, assume it contains $t = 0$.

$$g(t) = f(Z+tV) = \log\det\left[Z^{1/2}(I + tZ^{-1/2}VZ^{-1/2})Z^{1/2}\right]$$
$$= \log\det\left[(I + tZ^{-1/2}VZ^{-1/2})Z\right] = \log\det Z + \sum \log(1+t\lambda_i)$$

  Where λ_i are the eigenvalues of $Z^{-1/2}VZ^{-1/2}$. Taking derivatives of g, we find that the derivative is always ≤ 0. Thus, convexity.

- ***Use the epigraph***: Consider $f(\boldsymbol{x}, Y) = \boldsymbol{x}^\top Y^{-1}\boldsymbol{x}$. It is convex over $\mathbb{R}^n \times \mathbb{S}^n_{++}$.

$$\text{epi } f = \left\{(\boldsymbol{x}, Y, t) : \boldsymbol{x}^\top Y^{-1}\boldsymbol{x} \leq t, Y \succ 0\right\} = \left\{(\boldsymbol{x}, Y, t) : \begin{bmatrix} Y & \boldsymbol{x} \\ \boldsymbol{x}^\top & t \end{bmatrix} \succeq 0, Y \succ 0\right\}$$

  (We used Schur Complements). This is a set of LMIs, and therefore convex.

- ***Jensen's Inequality***: f convex $\Leftrightarrow$ $f\left(\mathbb{E}[\boldsymbol{x}]\right) \leq \mathbb{E}\left[f(\boldsymbol{x})\right]$, $\forall \boldsymbol{x}$ s.t. $\mathbb{P}(\boldsymbol{x} \in \text{dom } f) = 1$.

Operations that preserve convexity

- Non-negative weighed sum.
- The *perspective function* $g(\boldsymbol{x}, t) = tf(\boldsymbol{x}/t)$ $[t > 0]$ is convex if f is convex.

  Example: The perspective of the negative logarithm gives the relative entropy and Kullback-Leibler divergence. □

- The *pointwise maximum* $\sup_{\boldsymbol{y}\in\mathcal{Y}} f(\boldsymbol{x}, \boldsymbol{y})$ is an extended-value convex function, if $f(\cdot, \boldsymbol{y})$ is convex for each $\boldsymbol{y} \in \mathcal{Y}$ [note that we do *not* require *joint* convexity of f] This corresponds to intersections of epigraphs.

  Example: Let $f(\boldsymbol{x})$ be the sum of the r largest elements of $\boldsymbol{x}$. Then we can write $f(\boldsymbol{x})$ as the maximum of all the possible sums of r elements of $\boldsymbol{x}$. □

  Example: Support function $\sigma_{\mathcal{C}}(\boldsymbol{x}) = \sup\left\{\boldsymbol{x}\cdot\boldsymbol{y} : \boldsymbol{y} \in \mathcal{C}\right\}$. Convex. □

  Example: $\lambda_{\max}(X) = \sup\left\{\boldsymbol{x}^\top X\boldsymbol{x} : \left\|\boldsymbol{x}\right\| \leq 1\right\}$; family of linear functions of X. □

  Note: Every convex function $f : \mathbb{R}^n \to \mathbb{R}$ can be written as

$$f'(\boldsymbol{x}) = \sup\left\{g(\boldsymbol{x}) : g \text{ affine}, \boldsymbol{g}(\boldsymbol{z}) \leq f(\boldsymbol{z})\ \forall \boldsymbol{z}\right\}$$

Clearly, $f'(\boldsymbol{x}) \leq f(\boldsymbol{x})$. Furthermore, epi f is convex and so at any $(\boldsymbol{z}, t) \in \text{epi } f$, $\exists(\boldsymbol{\lambda}, \mu) \neq \boldsymbol{0}$ such that $\begin{bmatrix} \boldsymbol{\lambda} \\ \mu \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{x} - \boldsymbol{z} \\ f(\boldsymbol{x}) - t \end{bmatrix} \leq 0 \Rightarrow \boldsymbol{\lambda} \cdot (\boldsymbol{x} - \boldsymbol{z}) + \mu\left(f(\boldsymbol{x}) - t\right) \leq 0$ for all $\boldsymbol{x}$.

Now, (1) we must have $\mu \geq 0$, else as we take $t \to \infty$ violated *and* (2) we must have $\mu \neq 0$, else we get $\boldsymbol{\lambda} = \boldsymbol{0}$. Thus, can write $g(\boldsymbol{z}) = f(\boldsymbol{x}) + \frac{1}{\mu}\boldsymbol{\lambda} \cdot (\boldsymbol{x} - \boldsymbol{z}) \leq t$ Choosing a point on the boundary of the epigraph, $t = f(\boldsymbol{z})$ and so $g(\boldsymbol{z}) \leq t = f(\boldsymbol{z})$. As such, it's a global underestimator with $g(\boldsymbol{x}) = f(\boldsymbol{x})$. ☐

- The *minimum* $\inf_{y \in \mathcal{C}} f(\boldsymbol{x}, \boldsymbol{y})$ is convex provided it is $> -\infty$ and $f(\boldsymbol{x}, \boldsymbol{y})$ is **jointly** convex in ($\boldsymbol{x}$, $\boldsymbol{y}$) over $\mathbb{R}^n \times \mathbb{R}^m$ and $\mathcal{C} \subset \mathbb{R}^m$ is convex. This corresponds to the projective of a (convex) epigraph onto the subspace of $\boldsymbol{x}$.
- The *composition with an affine function* $f\left(A\boldsymbol{x} + \boldsymbol{b}\right)$ is convex, provided $f(\cdot)$ is convex.
- The general composition $f(\boldsymbol{x}) = h\left(g_1(\boldsymbol{x}), \cdots, g_k(\boldsymbol{x})\right)$ behaves as follows:

Result	*h*	*h*	*g*
Con**vex**	Con**vex**	Non**dec**reasing	Con**vex**
Con**vex**	Con**vex**	Non**inc**reasing	Con**cave**
Con**cave**	Con**cave**	Non**dec**reasing	Con**cave**
Con**cave**	Con**cave**	Non**inc**reasing	Con**vex**

These can be derived (for the case $k = 1$ and $\boldsymbol{x}$ scalar) by noting that the second derivative of f is given by

$$f''(x) = h''\left(g(x)\right)\left[g'(x)\right]^2 + h'\left(g(x)\right)g''(x)$$

Conjugate Functions

- $f^*(\boldsymbol{y}) = \sup_{\boldsymbol{x} \in \text{dom } f}\left(\boldsymbol{y} \cdot \boldsymbol{x} - f(\boldsymbol{x})\right)$. It's the maximum gap between the linear function $g(x) = yx$ and f. If f is differentiable, this occurs at a point at which $f'(\boldsymbol{x}) = \boldsymbol{y}$.

Basic examples:

***f*(*x*)**	$f^*(x)$	**Domain**
$ax + b$	$-b$	$\{a\}$
$-\log x$	$\log(1/y) - 1$	$-\mathbb{R}_{++}$
e^x	$y \log y - y$	$\mathbb{R}_+$

***f*(*x*)**	$f^*(x)$	**Domain**
$x \log x$	e^{y-1}	$\mathbb{R}$
$1/x$	$-2(-y)^{1/2}$	$-\mathbb{R}_+$

Example: $f(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^\top Q\boldsymbol{x}$ with $Q \in \mathbb{S}^n_{++}$. Then $\boldsymbol{y}\cdot\boldsymbol{x} - \frac{1}{2}\boldsymbol{x}^\top Q\boldsymbol{x}$ which is bounded above and maximized at $\boldsymbol{y} - Q\boldsymbol{x} = 0 \Rightarrow \boldsymbol{x} = Q^{-1}\boldsymbol{y}$, and $f^*(\boldsymbol{y}) = \frac{1}{2}\boldsymbol{y}^\top Q^{-1}\boldsymbol{y}$. □

Example: Let $I_{\mathcal{S}}(\boldsymbol{x}) = 0$ if $\boldsymbol{x} \in \mathcal{S}, \infty$ otherwise. Then the conjugate of this indicator is the support function $I^*_{\mathcal{S}}(\boldsymbol{y}) = \sup_{\boldsymbol{x}\in\mathcal{S}}(\boldsymbol{y}\cdot\boldsymbol{x}) = \sigma_{\mathcal{S}}(\boldsymbol{y})$. □

Example: If $f(\boldsymbol{x}) = \|\boldsymbol{x}\|$, then $f^*(\boldsymbol{y}) =$ Indicator f^n of $\left\{\boldsymbol{y} : \|\boldsymbol{y}\|_* \leq 1\right\}$. To see why; if $\|\boldsymbol{y}\|_* > 1$, then $\exists \boldsymbol{z}$ s.t. $\boldsymbol{y}\cdot\boldsymbol{z} > 1, \|\boldsymbol{z}\| \leq 1$. Take $\boldsymbol{x} = t\boldsymbol{z}$ and let $t \to \infty$. If $\|\boldsymbol{y}\|_* \leq 1$, then $\boldsymbol{x}\cdot\boldsymbol{y} \leq \|\boldsymbol{x}\|\|\boldsymbol{y}\|_* \leq \|\boldsymbol{x}\|$. Therefore, 0 maximizes $\boldsymbol{x}\cdot\boldsymbol{y} - \|\boldsymbol{x}\|$ maxed at 0. □

Example: $f(\boldsymbol{x}) = \log\left(\sum e^{x_i}\right)$. Differentiating $\boldsymbol{y}\cdot\boldsymbol{x} - f(\boldsymbol{x})$ and setting to 0, we get $y_i = e^{x_i} / \sum e^{x_i}$ and $f^*(\boldsymbol{x}) = \sum y_i \log y_i$ if $\boldsymbol{y} \geq \boldsymbol{0}, \boldsymbol{1}\cdot\boldsymbol{y} = 1$ and ∞ otherwise [this is valid even if some of the components of $\boldsymbol{y}$ are 0]. □

Example: Company uses resources $\boldsymbol{r}$ at price $\boldsymbol{p}$ prices produces revenue $S(\boldsymbol{r})$. The maximum profit that can be made from a given price is $M(\boldsymbol{p}) = \sup_r[S(\boldsymbol{r}) - \boldsymbol{p}\cdot\boldsymbol{r}] = (-S^*)(-p)$. □

- From the definition, we get $f(\boldsymbol{x}) + f^*(\boldsymbol{y}) \geq \boldsymbol{x}\cdot\boldsymbol{y}$.
- If f is differentiable, the maximizer of $\boldsymbol{y}\cdot\boldsymbol{x} - f(\boldsymbol{x})$ satisfies $\boldsymbol{y} = \nabla f(\boldsymbol{x})$. Thus, for any $\boldsymbol{z}$ for which $\boldsymbol{y} = \nabla f(\boldsymbol{z})$, $f^*(\boldsymbol{y}) = \boldsymbol{z}\cdot\nabla f(\boldsymbol{z}) - f(\boldsymbol{z})$
- The conjugate of $g(\boldsymbol{x}) = f(A\boldsymbol{x} + \boldsymbol{b})$ is $g^*(\boldsymbol{y}) = f^*(A^{-\top}\boldsymbol{y}) - \boldsymbol{b}^\top A^{-\top}\boldsymbol{y}$.
- The conjugate of the sum of functions is the sum of conjugates.
- Let f be a proper convex function. A vector $\boldsymbol{g}$ is a *subgradient* of f at $\boldsymbol{x}$ if $f(\boldsymbol{z}) \geq f(\boldsymbol{x}) + \boldsymbol{g}\cdot(\boldsymbol{x} - \boldsymbol{z})\ \forall \boldsymbol{z} \in \mathbb{R}^n$. $\partial f(\boldsymbol{x})$ is the set of all subgradients at $\boldsymbol{x}$. It is a closed and convex set. If $\boldsymbol{x}$ is in the interior of the domain of f, $\partial f(\boldsymbol{x})$ is non-empty, and if f is differentiable at $\boldsymbol{x}$, $\partial f(\boldsymbol{x}) = \{\nabla f(\boldsymbol{x})\}$.
  - If f is a proper convex function, then $\boldsymbol{x}\cdot\boldsymbol{y} = f(\boldsymbol{x}) + f^*(\boldsymbol{y}) \Leftrightarrow \boldsymbol{y} \in \partial f(\boldsymbol{x})$ [if f is closed, these are also equivalent to $\boldsymbol{x} \in \partial f^*(\boldsymbol{y})$.

    Proof: Follows directly from the definition of f^*. ■

- ***Chernoff Bounds & Large Deviations***
  - *Chernoff Bounds* – $\mathbb{1}_{\{X\geq t\}} \leq e^{\lambda(X-t)} \Rightarrow \mathbb{P}(X \geq t) \leq e^{-\lambda t}\mathbb{E}[e^{\lambda X}]\ \forall\lambda \geq 0$.

- Define the *cumulant generating function* $f(\lambda) = \log \mathbb{E}[e^{\lambda X}]$, and define $f_+(\lambda) = f(\lambda)$ if $\lambda \geq 0$ and ∞ otherwise.
- Making the Chernoff Bound as tight as possible, we get

$$\mathbb{P}(X \geq t) \leq \inf_{\lambda \geq 0} e^{-\lambda t}\mathbb{E}(e^{\lambda X}) = \exp\left(-\sup_{\lambda \geq 0}[\lambda t - f(\lambda)]\right) = \exp\left(-f_+^*[t]\right)$$

- Similarly, $\mathbb{P}(\boldsymbol{X} \in \mathcal{C}) \leq \mathbb{E}[e^{\boldsymbol{\lambda}\cdot\boldsymbol{X}+\mu}]$ provided $\boldsymbol{\lambda}\cdot\boldsymbol{z} + \mu \geq 0 \ \forall \boldsymbol{z} \in \mathcal{C}$. Then, defining $f(\lambda) = \log \mathbb{E}[e^{\boldsymbol{\lambda}\cdot\boldsymbol{X}}]$, we find that

$$\begin{aligned}\log \mathbb{P}(\boldsymbol{X} \in \mathcal{C}) &\leq \inf_{\boldsymbol{\lambda},\mu}\left\{\mu + f(\boldsymbol{\lambda}) : -\boldsymbol{\lambda}\cdot\boldsymbol{z} \leq \mu \ \forall \boldsymbol{z} \in \mathcal{C}\right\} \\ &= \inf_{\boldsymbol{\lambda}}\left\{\sup_{\boldsymbol{z}\in\mathcal{C}}(-\boldsymbol{\lambda}\cdot\boldsymbol{z}) + f(\boldsymbol{\lambda})\right\} \\ &= \inf_{\boldsymbol{\lambda}}\left\{S_{\mathcal{C}}(-\lambda) + f(\lambda)\right\} = -f_{\mathcal{C}}^*(0)\end{aligned}$$

- ***Example***: Let $\boldsymbol{X}$ be multivariate Gaussian: $\boldsymbol{X} \sim N(0, I)$. We then have $f(\lambda) = \frac{1}{2}\boldsymbol{\lambda}\cdot\boldsymbol{\lambda}$. Consider $\mathcal{C} = \{\boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}\}$. Now
  - By LP duality: $S_{\mathcal{C}}(-\lambda) = \sup_{\boldsymbol{z}\in\mathcal{C}}(-\boldsymbol{\lambda}\cdot\boldsymbol{z}) = \inf_{A^\top \boldsymbol{u} = -\boldsymbol{\lambda}, \boldsymbol{u} \geq \boldsymbol{0}} \boldsymbol{b}\cdot\boldsymbol{u}$
  - As such, $\log \mathbb{P}(\boldsymbol{X} \in \mathcal{C}) \leq \inf_{\boldsymbol{\lambda}} \boldsymbol{b}\cdot\boldsymbol{u} + \frac{1}{2}\boldsymbol{\lambda}\cdot\boldsymbol{\lambda}$ s.t. $\boldsymbol{u} \geq \boldsymbol{0}, A^\top\boldsymbol{u} + \boldsymbol{\lambda} = \boldsymbol{0}$
  - Eliminating $\boldsymbol{\lambda}$, we get $\log \mathbb{P}(\boldsymbol{X} \in \mathcal{C}) \leq \inf_{\boldsymbol{u} \geq \boldsymbol{0}} \boldsymbol{b}\cdot\boldsymbol{u} + \frac{1}{2}\boldsymbol{u}^\top AA^\top\boldsymbol{u}$
  - Using QP duality, we can write this program as $\sup_{\boldsymbol{\lambda} \geq \boldsymbol{0}} -\frac{1}{2}\left\|A^{-1}(\boldsymbol{\lambda} - \boldsymbol{b})\right\|_2^2$.
  - Intepreting $\boldsymbol{\lambda}$ as a slack variable, this becomes $\sup_{A\boldsymbol{x} \leq \boldsymbol{b}} -\frac{1}{2}\left\|\boldsymbol{x}\right\|_2^2$
  - As such, $\mathbb{P}(X \in \mathcal{C}) \leq \exp\left[-\frac{1}{2}\text{dist}(0, \mathcal{C})\right]$.

Chapter 4 – Convex Optimization Problems

- ***Terminology***

$$\begin{array}{lll} \min & f_0(x) & \\ \text{s.t.} & f_i(x) \leq 0 & i = 1, \cdots, m \\ & h_i(x) = 0 & i = 1, \cdots, p \end{array}$$

  - $\mathcal{D}$ is the feasible region of the problem
- ***Equivalent Problems***
  - Informal definition of equivalent problems: from the solution to one problem, a solution to the other is readily found, and vice versa.

- ***Change of variables*** – consider a one-one function $\phi : \mathbb{R}^n \to \mathbb{R}^n$ with $\mathcal{D} \subseteq \phi\left(\text{dom } \phi\right)$. Then replacing x by $\phi(x)$ leads to an equivalent problem.
  - ***Example (Linear-fractional programming)***: Consider

    $$\min \tfrac{\boldsymbol{c}\cdot\boldsymbol{x}+\boldsymbol{d}}{\boldsymbol{e}\cdot\boldsymbol{x}+\boldsymbol{f}} \text{ s.t. } A\boldsymbol{x} = \boldsymbol{b}, G\boldsymbol{x} \le \boldsymbol{h}$$

    Simply write the objective as $\boldsymbol{c} \cdot \left(\frac{x}{\boldsymbol{e}\cdot\boldsymbol{x}+\boldsymbol{f}}\right) + \boldsymbol{d}\left(\frac{1}{\boldsymbol{e}\cdot\boldsymbol{x}+\boldsymbol{f}}\right)$, and min $\boldsymbol{c}\cdot\boldsymbol{y} + \boldsymbol{d}z$.
- ***Transformation of objectives & Constraints*** – suppose:
  - $\psi_0 : \mathbb{R} \to \mathbb{R}$ is monotone increasing.
  - $\psi_1, \cdots, \psi_m : \mathbb{R} \to \mathbb{R}$ satisfy $\psi_i(u) \le u \Leftrightarrow u \le 0$
  - $\psi_{m+1}, \cdots, \psi_{m+p} : \mathbb{R} \to \mathbb{R}$ satisfy $\psi_i(u) = 0 \Leftrightarrow u = 0$

  Then composing f and h with these functions leads to the same problem.
- ***Eliminating equality constraints*** – say we find a function $\phi : \mathbb{R}^k \to \mathbb{R}^n$ such that x satisfies the equality constraints if and only if it can be written as $x = \phi(z)$. Then we can eliminate the equality constraints and optimize $f_0(\phi(z))$ s.t. $f_i(\phi(z)) \le 0$ over z. For example, the equality constraint $A\boldsymbol{x} = \boldsymbol{b}$ with $A \in \mathbb{R}^{m\times n}$ (with solution $\boldsymbol{x}_0$) can be replaced by $\boldsymbol{x} = F\boldsymbol{z} + \boldsymbol{x}_0$, with $F \in \mathbb{R}^{n\times[n-\text{rank } A]}$. This preserves convexity, since this is an affine transformation.

  Optimizing over some variables – it is possible to optimize over each variable one-by-one; this is especially true when some constraints involve only a subset of the variables. For example, take $\min_{\boldsymbol{x}:f_i(\boldsymbol{x}_1)\le 0} \boldsymbol{x}^\top P\boldsymbol{x}$. Minimizing over x_2 only gives an objective of $\boldsymbol{x}_1^\top S\boldsymbol{x}_1$. Thus, the problem is equivalent to $\min_{\boldsymbol{x}_1:f_i(\boldsymbol{x}_1)\le 0} \boldsymbol{x}_1^\top S\boldsymbol{x}_1$.
- ***Epigraph form*** – the problem above is equivalent to minimizing t subject to $f_0(x) - t \le 0$. Particularly useful for minimax problem.

  Example: $\min_{\boldsymbol{x}} \max_{1\le i\le r} \left\|\boldsymbol{x} - \boldsymbol{y}_i\right\| = \min_{\boldsymbol{x},t} t \text{ s.t. } t \ge \left\|\boldsymbol{x} - \boldsymbol{y}_i\right\|^2 \ \forall i$ – a QCQP. □
- ***Implicit & Explicit constraints*** – if the objective function has a restricted domain, it can often be unrestricted by adding additional constraints instead, and vice versa.

- ***Convex problems***

A convex optimization problem is

$$\begin{array}{llll} \min & f_0(\boldsymbol{x}) & & \leftarrow \text{convex} \\ \text{s.t.} & f_i(\boldsymbol{x}) \le 0 & i = 1, \cdots, m & \leftarrow \text{convex} \\ & a_i^\top x = b_i & i = 1, \cdots, p & \leftarrow \text{afine} \end{array}$$

If f_0 is *quasiconvex*, the problem is a *quasiconvex optimization problem*. In either case, the ε-suboptimal sets are convex. An optimality condition for $\boldsymbol{x}$ is

$$\nabla f_0(\boldsymbol{x}) \cdot (\boldsymbol{y} - \boldsymbol{x}) \ge 0 \qquad \forall \boldsymbol{y} \text{ feasible}$$

Geometrically, $\nabla f_0(\boldsymbol{x})$ is a supporting hyperplane to the feasible set at $\boldsymbol{x}$. Alternatively, anywhere feasible we try to move from $\boldsymbol{x}$ yields an increase in objective.

Proof: For a convex function, $f_0(\boldsymbol{y}) \ge f_0(\boldsymbol{x}) + \nabla f_0(\boldsymbol{x}) \cdot (\boldsymbol{y} - \boldsymbol{x})$. If the optimality condition is met, $f_0(\boldsymbol{y}) \ge f_0(\boldsymbol{x}) \ \forall \boldsymbol{y}$. *Conversely*, suppose $\exists \boldsymbol{y}$ s.t. $\nabla f_0(\boldsymbol{x}) \cdot (\boldsymbol{y} - \boldsymbol{x}) < 0$. Then consider $\boldsymbol{z}(t) = t\boldsymbol{y} + (1-t)\boldsymbol{x}, t \in [0,1]$. This is feasible, but $\frac{\mathrm{d}}{\mathrm{d}t} f_0(\boldsymbol{z}(t))\big|_{t=0} = \nabla f_0(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) < 0$. So close to 0, we can decrease f_0 by moving away from f. ■

Examples:

- For unconstrained problems, $\nabla f_0(\boldsymbol{x}) = 0$
- For problems with $A\boldsymbol{x} = \boldsymbol{b}$, every solution can be written as $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{v}, \boldsymbol{v} \in \mathcal{N}(A)$ So $\nabla f_0(\boldsymbol{x}) \cdot \boldsymbol{v} \ge 0 \ \forall \boldsymbol{v} \in \mathcal{N}(A)$, but this is a nullspace, so $\nabla f_0(\boldsymbol{x}) \cdot \boldsymbol{v} = 0 \ \forall \boldsymbol{v} \in \mathcal{N}(A)$ In other words, $\nabla f_0(\boldsymbol{x}) \in \mathcal{N}(A)^\perp = \mathcal{R}(A^\top)$, ie: $\exists \boldsymbol{\nu}$ s.t. $\nabla f_0(\boldsymbol{x}) + A^\top \boldsymbol{\nu} = \boldsymbol{0}$.
- For problems with $\boldsymbol{x} \succeq \boldsymbol{0}$, need $\nabla f_0(\boldsymbol{x}) \ge \boldsymbol{0}$, else $\nabla f_0(\boldsymbol{x}) \cdot (\boldsymbol{y} - \boldsymbol{x})$ unbounded below. So reduces to $\nabla f_0(\boldsymbol{x}) \cdot \boldsymbol{x} \ge 0$. This gives complementary slackness…

Examples of convex optimization problems.

- ***Linear programs (LP)***
  - ***Example (Chebyshev Center)***: Consider the problem of finding the largest ball $\mathcal{B} = \left\{ \boldsymbol{x}_c + \boldsymbol{u} \mid \|\boldsymbol{u}\| \le r \right\}$ that lies in a polyhedron. Require that $\mathcal{B}$ lies on one side of $\boldsymbol{a} \cdot \boldsymbol{x} \le \boldsymbol{b}$ is equivalent to requiring:

    $$\sup_{\|\boldsymbol{u}\| \le r} \left\{ \boldsymbol{a} \cdot \boldsymbol{x}_c + \boldsymbol{a} \cdot \boldsymbol{u} \right\} \le \boldsymbol{b} \Rightarrow \boldsymbol{a} \cdot \boldsymbol{x}_c + r\|\boldsymbol{a}\| \le \boldsymbol{b} \qquad \square$$

  - ***Example***: $\min_{\boldsymbol{x}} \max_i \left(\boldsymbol{a}_i \cdot \boldsymbol{x} + \boldsymbol{b}_i \right)$ can be linearised using the epigraph. □
- ***Linear-fractional programs***

$$\min_{x: e\cdot x + f > 0} \quad \frac{\boldsymbol{c}\cdot\boldsymbol{x}+\boldsymbol{d}}{\boldsymbol{e}\cdot\boldsymbol{x}+\boldsymbol{f}} \qquad \text{s.t. } \boldsymbol{x}\in\mathcal{P}=\left\{\boldsymbol{x}: G\boldsymbol{x}\le\boldsymbol{h}, A\boldsymbol{x}=\boldsymbol{b}\right\}$$

This is a quasiconvex program, which can be transformed into

$$\begin{aligned}\min_{\boldsymbol{y},z} \quad & \boldsymbol{c}\cdot\boldsymbol{y}+\boldsymbol{d}z \\ \text{s.t.} \quad & G\boldsymbol{y}-\boldsymbol{h}z\le 0, A\boldsymbol{y}-\boldsymbol{b}z=0, \boldsymbol{e}\cdot\boldsymbol{y}+\boldsymbol{f}z=1, z\ge 0\end{aligned}$$

Proof: Note that $z=\frac{1}{e\cdot x+f}, \boldsymbol{y}=\boldsymbol{x}z$ is feasible for the transformed problem with the same objective. Similarly, $\boldsymbol{x}=\boldsymbol{y}/z$ is feasible for the original problem, if $z\neq 0$. Thus, the original and modified objectives are both $\ge$ and $\le$ each other. If $z=0$ and $\boldsymbol{x}_0$ is feasible for the original problem, then $\boldsymbol{x}=\boldsymbol{x}_0+t\boldsymbol{y}$ is optimal for the original problem; taking $t\to\infty$ allows us to gets the two objectives arbitrarily close to each other. ☐

- ***Quadratic programs (QP)***

$$\min_{\boldsymbol{x}} \quad \tfrac{1}{2}\boldsymbol{x}^\top P\boldsymbol{x}+\boldsymbol{q}\cdot\boldsymbol{x}+\boldsymbol{r} \qquad \text{s.t. } \boldsymbol{x}\in\mathcal{P}=\left\{G\boldsymbol{x}\le\boldsymbol{h}, A\boldsymbol{x}=\boldsymbol{b}\right\}$$

Where $P\succeq 0$.

  - ***Example***: Distance between polyhedra $\min\left\|\boldsymbol{x}_1-\boldsymbol{x}_2\right\|_2^2$ s.t. $\boldsymbol{x}_1\in\mathcal{P}_1, \boldsymbol{x}_2\in\mathcal{P}_2$.
  - ***Example***: Consider an LP minimizing $\boldsymbol{c}\cdot\boldsymbol{x}$ where $\mathbb{E}(\boldsymbol{x})=\overline{\boldsymbol{x}}$, $\mathbb{C}\text{ov}(\boldsymbol{x})=\Sigma$. Then $\mathbb{V}\text{ar}(\boldsymbol{c}\cdot\boldsymbol{x})=\boldsymbol{x}^\top\Sigma\boldsymbol{x}$. Can minimize $\boldsymbol{c}\cdot\boldsymbol{x}+\gamma\boldsymbol{x}^\top\Sigma\boldsymbol{x}$. ☐
  - ***Example (portfolio problems)***: p_i is relative price change of asset i (change/start) and x_i is amount of asset bought. Mean return is $\boldsymbol{p}\cdot\boldsymbol{x}$, variance is $\boldsymbol{x}^\top\Sigma\boldsymbol{x}$. Minimize variance subject to given mean return. Budget constraint $\boldsymbol{1}\cdot\boldsymbol{x}=1$. ***Short sales***: Can set $\boldsymbol{x}=\boldsymbol{x}^{\text{long}}-\boldsymbol{x}^{\text{short}}$ and require $\boldsymbol{1}\cdot\boldsymbol{x}^{\text{short}}\le\eta\boldsymbol{1}\cdot\boldsymbol{x}^{\text{long}}$. ***Transaction costs***: Can set $\boldsymbol{x}=\boldsymbol{x}^{\text{init}}+\boldsymbol{u}^{\text{buy}}-\boldsymbol{u}^{\text{sell}}$ and constraint $\left(1-f^{\text{sell}}\right)\boldsymbol{1}\cdot\boldsymbol{u}^{\text{sell}}=\left(1+f^{\text{buy}}\right)\boldsymbol{1}\cdot\boldsymbol{u}^{\text{buy}}$. ☐

- ***Quadratically constrained quadratic programs (QCQP)***

$$\min_{\boldsymbol{x}} \quad \tfrac{1}{2}\boldsymbol{x}^\top P_0\boldsymbol{x}+\boldsymbol{q}_0\cdot\boldsymbol{x}+\boldsymbol{r}_0 \qquad \text{s.t. } \tfrac{1}{2}\boldsymbol{x}^\top P_i\boldsymbol{x}+\boldsymbol{q}\cdot\boldsymbol{x}+\boldsymbol{r}_i\le\boldsymbol{0}, \boldsymbol{x}\in\mathcal{P}$$

Where the P_i are positive definite. In this case, we are maximizing a quadratic function over an intersection of ellipses and polyhedra.

- ***Second-order cone program (SOCP)***

$$\min_{\boldsymbol{x}} \boldsymbol{f}^\top\boldsymbol{x} \text{ s.t. } \left\|A^i\boldsymbol{x}+\boldsymbol{b}^i\right\|_2\le\boldsymbol{c}\cdot\boldsymbol{x}+\boldsymbol{d} \quad i=,1,\cdots,m, \ \boldsymbol{x}\in\mathcal{P}$$

If $\boldsymbol{c}_i = \boldsymbol{0}\ \forall i$, this reduces to a QCQP (square both sides to see). Note that the direction of the inequality is important!

- ***Example (quadratic constraint)***:

$$\boldsymbol{x}^\top A^\top A\boldsymbol{x} + \boldsymbol{b}\cdot\boldsymbol{x} + \boldsymbol{c} \le 0 \Leftrightarrow \left\| \begin{matrix} \frac{1}{2}(1+\boldsymbol{b}\cdot\boldsymbol{x}+\boldsymbol{c}) \\ A\boldsymbol{x} \end{matrix} \right\|_2 \le \frac{1}{2}\left(1-\boldsymbol{b}\cdot\boldsymbol{x}-\boldsymbol{c}\right)$$

- ***Example (robust LP)***: Consider $\min_{\boldsymbol{x}} \boldsymbol{c}\cdot\boldsymbol{x}$ s.t. $\boldsymbol{a}^i\cdot\boldsymbol{x}\le \boldsymbol{b}^i$ for all $\boldsymbol{a}^i \in \mathcal{E}_i = \left\{\bar{\boldsymbol{a}}^i + P^i\boldsymbol{u} \mid \left\|\boldsymbol{u}\right\| \le 1\right\}$ (ie: ellipsoids). Can write constraint as $\bar{\boldsymbol{a}}^i\cdot\boldsymbol{x} + \sup\left\{\boldsymbol{u}^\top P^{i,\top}\boldsymbol{x} \mid \left\|\boldsymbol{u}\right\|\le 1\right\} \le \boldsymbol{b}^i \Leftrightarrow \bar{\boldsymbol{a}}^i\cdot\boldsymbol{x} + \left\|P^{i,\top}\boldsymbol{x}\right\| \le \boldsymbol{b}^i$, an SOCP. Norm terms are *regularization terms*; prevents $\boldsymbol{b}$ from being high in directions where uncertainty in $\boldsymbol{a}$ is high. □

  Note: If $\mathcal{E} = \mathcal{P} = \left\{\boldsymbol{x} : A\boldsymbol{x} = \boldsymbol{f}\right\}$ is a polyhedron with k vertices, then $\sup\left\{\boldsymbol{a}\cdot\boldsymbol{x} : \boldsymbol{x}\in\mathcal{P}\right\}$ must occur at a vertex, so 1 inequality just turns into k inequalities. Alternatively, $\sup_{A\boldsymbol{x}\le \boldsymbol{f}}\left\{\boldsymbol{a}\cdot\boldsymbol{x}\right\} = \inf_{A^\top \boldsymbol{z}=\boldsymbol{a},\boldsymbol{z}\ge\boldsymbol{0}}\left\{\boldsymbol{f}\cdot\boldsymbol{z}\right\}$ (strong duality) and so write $\boldsymbol{a}\cdot\boldsymbol{x}\le\boldsymbol{b} \Leftrightarrow \boldsymbol{f}\cdot\boldsymbol{z}\le\boldsymbol{b}, A^\top\boldsymbol{z} = \boldsymbol{a}, \boldsymbol{z}\ge 0$. □

- ***Example (uncertain LP)***: Say $\boldsymbol{a}^i \sim N(\bar{\boldsymbol{a}}^i, \Sigma_i)$. Can ask for $\mathbb{P}(\boldsymbol{a}^i\cdot\boldsymbol{x}\le\boldsymbol{b}^i) \ge \eta$. Note that $\mathbb{SD}(\boldsymbol{a}^i\cdot\boldsymbol{x}) = \sqrt{\boldsymbol{x}^\top\Sigma_i\boldsymbol{x}} = \left\|\Sigma_i^{1/2}\boldsymbol{x}\right\|_2$. Can then express probability constraint as SOCP. □

- ***Semidefinite programs*** (SDP)

$$\min_{\boldsymbol{x}} \boldsymbol{c}\cdot\boldsymbol{x} \text{ s.t. } \boldsymbol{x}_1 F_1 + \cdots + \boldsymbol{x}_n F_n + G \preceq 0, \boldsymbol{x}\in\mathcal{P}$$

Multiple LMIs are easily dealt with by forming a large block diagonal LMI from the individual LMIs. If the matrices F are diagonal, this is an LP.

- ***Example (bounds on eigenvalues)***: Consider min $\left\|A\right\|_2$ ($\left\|A\right\|_2$ is max singular value). Note $\left\|A\right\|_2 \le s \Leftrightarrow A^\top A \preceq s^2 I$. Using Schur complements,

$$A^\top A - s^2 I \preceq 0 \Leftrightarrow \begin{bmatrix} sI & A \\ A^\top & sI \end{bmatrix} \succeq 0$$

  (Clearly, $sI \succ 0$ and the Schur complement is $S = sI - \frac{AA^\top}{s}$). Thus, simplify minimize s subject to the constraint above. The original LMI was quadratic – Schur complements allowed us to make it linear. [Similarly, we can bound the *lowest* eigenvalue: $\lambda_{\min}(A) \geq s \Leftrightarrow A \succeq sI$] ☐

- ***Example (portfolio optimization)***: Say we know $L_{ij} \leq \Sigma_{ij} \leq U_{ij}$. Given a portfolio $\boldsymbol{x}$, can maximize $\boldsymbol{x}^\top \Sigma \boldsymbol{x}$ s.t. that constraint and $\Sigma \succeq 0$ to get worst-case variance. We can add additional convex constraints
  - ***Known portfolio variances***: $\boldsymbol{u}_k^\top \Sigma \boldsymbol{u}_k = \sigma_k^2$
  - ***Estimation error***: If we estimate $\Sigma = \hat{\Sigma}$ but within an ellipsoidal confidence interval, we have $C\left(\Sigma - \hat{\Sigma}\right) \leq \alpha$, where $C(\cdot)$ is some positive definite quadratic form.
  - ***Factor models***: Say $p = Fz + d$, where z are random *factors* and d represents additional randomness. We then have $\Sigma = F\Sigma_{\text{factor}}F^\top + D$, and we can constraint each individually.
  - ***Correlation coefficients***: $\rho_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\Sigma_{jj}}}$. In a case where we know the volatilities exactly, constraints on ρ_{ij} are linear… ☐
- ***Example (expressing QCQP and SOCP as SDP)***: Using Schur complements, we can make these non-linear constraints linear

$$\tfrac{1}{2}\boldsymbol{x}^\top P\boldsymbol{x} + \boldsymbol{g}\cdot\boldsymbol{x} + \boldsymbol{h} \leq 0 \Leftrightarrow \begin{bmatrix} I & P^{1/2}\boldsymbol{x} \\ (P^{1/2}\boldsymbol{x})^\top & -\boldsymbol{g}\cdot\boldsymbol{x} - \boldsymbol{h} \end{bmatrix} \succeq 0$$

$$\left\|F\boldsymbol{x} + \boldsymbol{q}\right\| \leq \boldsymbol{g}\cdot\boldsymbol{x} + \boldsymbol{h} \Leftrightarrow \begin{bmatrix} (\boldsymbol{g}\cdot\boldsymbol{x} + \boldsymbol{h})I & F\boldsymbol{x} + \boldsymbol{q} \\ (F\boldsymbol{x} + \boldsymbol{q})^\top & \boldsymbol{g}\cdot\boldsymbol{x} + \boldsymbol{h} \end{bmatrix} \succeq 0$$

- ***Geometric Programming***
  - A function $f(\boldsymbol{x}) = \sum_{k=1}^{K} c_k x_1^{a_{1k}} x_2^{a_{2k}} \cdots x_n^{a_{nk}}$, with $c_k > 0$ and $a_i \in \mathbb{R}$ is a ***posynomial*** (closed under $+, \times$). $K = 1$ gives a ***monomial*** (closed under $\times, \div$).
    - Posynomial $\times$ Monomial $=$ Posynomial
    - Posynomial $\div$ Monomial $=$ Posynomial
  - A ***geometric program*** is of the form

$$\min f_0(\boldsymbol{x}) \text{ s.t. } f_i(\boldsymbol{x}) \leq 1, h_i(\boldsymbol{x}) = 1, \boldsymbol{x} > \boldsymbol{0}$$

f are posynomials, h are monomials. Can deal with $f(\boldsymbol{x}) \le g(\boldsymbol{x})$ and $h(\boldsymbol{x}) = g(\boldsymbol{x})$ by dividing. Can maximize by minimizing inverse (also posynomial).

- ***Example***: $\max x_1 x_2 x_3$ s.t. $x_1 x_2 + x_2 x_3 + x_1 x_3 \le c / 2, \boldsymbol{x} > \boldsymbol{0}$ [min volume box] can be written as $\min x_1^{-1} x_2^{-1} x_3^{-1}$ s.t. $(2x_1 x_2 / c) + (2x_2 x_3 / c) + (2x_1 x_3 / c) \le 1, \boldsymbol{x} > \boldsymbol{0}$. ☐
- To make convex, substitute $y_i = \log x_i \Rightarrow x_i = e^{y_i}$. Feed in and then take logs of objective and constraints. Result is convex.

- ***Existence of Solutions***
  - ***Theorem*** (*Weierstrass*): Consider the problem $\min f(\boldsymbol{x})$ s.t. $\boldsymbol{x} \in \mathcal{C} \subset \mathbb{R}^n$. Then, if $\mathcal{C}$ is non-empty, f is lower semicontinuous over $\mathcal{C}$ and either (1) $\mathcal{C}$ is compact (2) $\mathcal{C}$ is closed, and f is coercive (3) There exists a scalar γ such that the level set $\mathcal{C}(\gamma) = \left\{\boldsymbol{x} \in \mathcal{C} : f(\boldsymbol{x}) \le \gamma\right\}$ is nonempty and compact, **then** the set of optimal minimizing solutions of f is non-empty and compact.

    Proof: Let f^* be the optimal objective, and $\gamma_k \downarrow f^*$. Then the set of optimal solutions is $\bigcap_{k=1}^{\infty} \mathcal{C}(\gamma_k)$. If (3) is true, this is an intersection of nested non-empty compact sets – it is therefore non-empty and compact. $\boxed{1 \to 3}$ For $\gamma > f^*$, $\mathcal{C}(\gamma)$ must be non-empty, and by semi-continuity of f, it is closed. Since $\mathcal{C}$ is compact, this closed subset is also compact. $\boxed{2 \to 3}$ $\mathcal{C}(\gamma) = f^{-1}\left((-\infty, \gamma] \cap \mathcal{C}\right)$ since $\mathcal{C}$ is closed, the intersection is closed and so is the inverse by semi-continuity of f. Since f is coercive, $\mathcal{C}(\gamma)$ is also bounded. Thus, $\mathcal{C}(\gamma)$ is compact. ■
  - ***Example***: Consider $\min \frac{1}{2}\boldsymbol{x}^\top P \boldsymbol{x} - \boldsymbol{b} \cdot \boldsymbol{x}, \boldsymbol{x} \in \mathbb{R}$. If λ is the smallest eigenvalue of P, then we can say that $\frac{1}{2}\boldsymbol{x}^\top P \boldsymbol{x} - \boldsymbol{b} \cdot \boldsymbol{x} \ge \frac{1}{2}\lambda \left\|\boldsymbol{x}\right\|^2 - \left\|\boldsymbol{b}\right\|\left\|\boldsymbol{x}\right\|$. This is coercive if $\lambda > 0$. Thus, a solution exists if $P \succ 0$. ☐

Chapter 5 – Duality

- ***The Lagrangian Function***
  - $\boxed{\min\ f_0(\boldsymbol{x}) \text{ s.t. } \boldsymbol{f}(\boldsymbol{x}) \le \boldsymbol{0}, \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}}$. We let $\mathcal{L}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = f_0(\boldsymbol{x}) + \boldsymbol{\lambda} \cdot \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{\nu} \cdot \boldsymbol{h}(\boldsymbol{x})$ be the Lagrangian, and $\boldsymbol{g}(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\boldsymbol{x}} \mathcal{L}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\nu})$. This clearly *underestimates* the optimal value, because everywhere in the feasible region, $\mathcal{L}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \le f_0(\boldsymbol{x})$.

- Writing the original program as $\min f_0(\boldsymbol{x}) + \sum_{i=1}^m I_-\left[f_i(\boldsymbol{x})\right] + \sum_{i=1}^p I_0\left[h_i(\boldsymbol{x})\right]$, where I_- and I_0 are indicator functions for the negative orthant and $\{0\}$, we see the Lagrangian replaces the indicators ("hard walls") by "soft walls".
- ***Example***: $\min \ \boldsymbol{x}\cdot\boldsymbol{x}$ s.t. $A\boldsymbol{x} = \boldsymbol{b}$. $\mathcal{L} = \boldsymbol{x}\cdot\boldsymbol{x} + \boldsymbol{\lambda}\cdot(A\boldsymbol{x}-\boldsymbol{b})$. Differentiating, set to 0, minimum is at $2\boldsymbol{x} + A^\top\boldsymbol{\lambda} = 0$, so $g(\boldsymbol{\lambda}) = \mathcal{L}\left(-\frac{1}{2}A^\top\boldsymbol{\lambda}, \boldsymbol{\lambda}\right)$. □

  Example: $\min \ \boldsymbol{x}^\top W\boldsymbol{x}$ s.t. $x_i^2 = 1$. Involves partitioning the x_i into either a "+1" group or a "–1" group; $\pm W_{ij}$ is the cost of having x_i and x_j in the same/different partitions. $\mathcal{L}(\boldsymbol{x},\boldsymbol{\nu}) = \boldsymbol{x}^\top W\boldsymbol{x} + \boldsymbol{\nu}\cdot(\boldsymbol{x}^2 - \boldsymbol{1}) = \boldsymbol{x}^\top\left[W + \operatorname{diag}(\boldsymbol{\nu})\right]\boldsymbol{x} - \boldsymbol{1}\cdot\boldsymbol{\nu}$. Minimizing over $\boldsymbol{x}$, we find that $g(\boldsymbol{\nu}) = -\boldsymbol{1}\cdot\boldsymbol{\nu}$ if $\left[W + \operatorname{diag}(\boldsymbol{\nu})\right] \succeq 0$ and $-\infty$ o.w. Can use to find bound – for example, using $\boldsymbol{\nu} = -\lambda_{\min}(W)\boldsymbol{1}$, feasible because $W - \lambda_{\min} I \succeq 0$. □

- ***The Lagrangian & Convex Conjugates***
  - Consider $\min f(\boldsymbol{x})$ s.t. $A\boldsymbol{x} \le \boldsymbol{b}, C\boldsymbol{x} = \boldsymbol{d}$. The dual function is

$$\begin{aligned} g(\boldsymbol{\lambda},\boldsymbol{\nu}) &= \inf_{\boldsymbol{x}}\left[f(\boldsymbol{x}) + \boldsymbol{\lambda}\cdot(A\boldsymbol{x}-\boldsymbol{b}) + \boldsymbol{\nu}\cdot(C\boldsymbol{x}-\boldsymbol{d})\right] \\ &= -\boldsymbol{\lambda}\cdot\boldsymbol{b} - \boldsymbol{\nu}\cdot\boldsymbol{d} + \inf_{\boldsymbol{x}}\left[f(\boldsymbol{x}) + (\boldsymbol{\lambda}^\top A + \boldsymbol{\nu}^\top C)\boldsymbol{x}\right] \\ &= -\boldsymbol{\lambda}\cdot\boldsymbol{b} - \boldsymbol{\nu}\cdot\boldsymbol{d} - \sup_{\boldsymbol{x}}\left[(-A^\top\boldsymbol{\lambda} - C^\top\boldsymbol{\nu})^\top\boldsymbol{x} - f(\boldsymbol{x})\right] \\ &= -\boldsymbol{\lambda}\cdot\boldsymbol{b} - \boldsymbol{\nu}\cdot\boldsymbol{d} - f^*(-A^\top\boldsymbol{\lambda} - C^\top\boldsymbol{\nu}) \end{aligned}$$

  Example: $f(\boldsymbol{x}) = \left\|\boldsymbol{x}\right\|$ and only equality constraints. The conjugate of a norm is the indicator of the unit ball in its dual norm, so $g(\boldsymbol{\nu}) = \begin{cases} -\boldsymbol{\nu}\cdot\boldsymbol{d} & \left\|C^\top\boldsymbol{\nu}\right\|_* \le 1 \\ -\infty & \text{otherwise} \end{cases}$. □

  Example: Min entropy: $\min \sum x_i \log x_i$ s.t. $A\boldsymbol{x} \le \boldsymbol{b}, \boldsymbol{1}\cdot\boldsymbol{x} = 1$. The conjugate of $x \log x$ is e^{v-1}, and so $g(\boldsymbol{\lambda},\nu) = -\boldsymbol{\lambda}\cdot\boldsymbol{b} - \nu - \sum \exp\left[-(A^\top\boldsymbol{\lambda})_i - \nu - 1\right]$. □

- ***The Lagrange Dual Problem***
  - The dual problem is $\max g(\boldsymbol{\lambda},\boldsymbol{\nu})$ s.t. $\boldsymbol{\lambda} \ge \boldsymbol{0}$, with domain $\left\{(\boldsymbol{\lambda},\boldsymbol{\nu}) : g(\boldsymbol{\lambda},\boldsymbol{\nu}) > -\infty\right\}$.
  - *Weak duality* implies that $d^* \le p^*$. If the primal is unbounded below, the dual is infeasible. If the dual is unbounded above, the primal is infeasible.
  - The dual is always convex and can be used to find a good lower bound.
  - *Strong duality* holds under Slater's Conditions; the problem is convex, and there exists a point such that every inequality constraint is strictly satisfied. If the constraints are linear, only feasibility is needed, not *strict* feasibility.

- *Example*: QCQP: $\min \frac{1}{2}\boldsymbol{x}^\top P_0 \boldsymbol{x} + \boldsymbol{q}_0 \cdot \boldsymbol{x} + \boldsymbol{r}_0$ s.t. $\frac{1}{2}\boldsymbol{x}^\top P_i \boldsymbol{x} + \boldsymbol{q}_i \cdot \boldsymbol{x} + \boldsymbol{r}_i \leq 0$, where $P_0 \succ 0, P_i \succeq 0$. The Lagrangian is

$$\mathcal{L}(\boldsymbol{x},\boldsymbol{\lambda}) = \tfrac{1}{2}\boldsymbol{x}^\top \overbrace{\left(P_0 + \sum \lambda_i P_i\right)}^{P(\boldsymbol{\lambda})}\boldsymbol{x} + \overbrace{\left(\boldsymbol{q}_0 + \sum \lambda_i \boldsymbol{q}_i\right)}^{\boldsymbol{q}(\lambda)}\cdot \boldsymbol{x} + \overbrace{\left(\boldsymbol{r}_0 + \boldsymbol{\lambda}\cdot \boldsymbol{r}_i\right)}^{\boldsymbol{r}(\lambda)}$$

  $\boldsymbol{\lambda} \geq \boldsymbol{0}$ and so $P(\boldsymbol{\lambda}) \succ 0$. Differentiating and setting to 0, we find the dual problem is $\max g(\boldsymbol{\lambda}) = -\frac{1}{2}\boldsymbol{q}(\boldsymbol{\lambda})^\top \mathrm{P}(\boldsymbol{\lambda})^{-1}\boldsymbol{q}(\boldsymbol{\lambda}) + \boldsymbol{r}(\boldsymbol{\lambda})$ s.t. $\boldsymbol{\lambda} \succeq \boldsymbol{0}$. $p^* = d^*$ if we have strict feasibility. □

- *Example*: Min entropy (above) has dual $\max_{\boldsymbol{\lambda}\geq 0,\nu}\left(-\boldsymbol{\lambda}\cdot\boldsymbol{b} - \nu - e^{-\nu-1}\sum e^{-(A^\top\boldsymbol{\lambda})_i}\right)$. Optimizing over ν, we get $\max_{\boldsymbol{\lambda}\geq 0}\left(-\boldsymbol{\lambda}\cdot\boldsymbol{b} - \sum e^{-(A^\top\boldsymbol{\lambda})_i} + 1\right)$, a GP. □

Geometric Interpretation/Proof of Strong Duality

- Consider $\mathcal{G} = \left\{(\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{h}(\boldsymbol{x}), f_0(\boldsymbol{x})\right\}$. Then $p^* = \inf\left\{t : (\boldsymbol{u},\boldsymbol{v},t)\in\mathcal{G}, \boldsymbol{u}\leq\boldsymbol{0}, \boldsymbol{v}=\boldsymbol{0}\right\}$, and $g(\boldsymbol{\lambda},\boldsymbol{\nu}) = \inf\left\{(\boldsymbol{\lambda},\boldsymbol{\nu},1)\cdot(\boldsymbol{u},\boldsymbol{v},t) : (\boldsymbol{u},\boldsymbol{v},t)\in\mathcal{G}\right\}$. To assuming the infemum exists, the inequality $(\boldsymbol{\lambda},\boldsymbol{\nu},1)\cdot(\boldsymbol{u},\boldsymbol{v},t)\geq g(\boldsymbol{\lambda},\boldsymbol{\nu})\ \forall(\boldsymbol{u},\boldsymbol{v},t)\in\mathcal{G}$ defines a *supporting hyperplane*. Looking at the expression for p^* and $g(\boldsymbol{\lambda},\boldsymbol{\nu})$, it is clear that $p^* \geq g(\boldsymbol{\lambda},\boldsymbol{\nu})$ when $\boldsymbol{\lambda}\geq\boldsymbol{0}$, because it involves more constraints.
- Let $\mathcal{A} = \mathcal{G} + \left(\mathbb{R}^m_+, \{\boldsymbol{0}\}, \mathbb{R}_+\right)$ be an "epigraph" of $\mathcal{G}$ which extends it "up" in the objective & inequalities. This then allows us to write $p^* = \inf\left\{t : (\boldsymbol{0},\boldsymbol{0},t)\in\mathcal{A}\right\}$, and, for $\boldsymbol{\lambda}\geq\boldsymbol{0}$, $g(\boldsymbol{\lambda},\boldsymbol{\nu}) = \inf\left\{(\boldsymbol{\lambda},\boldsymbol{\nu},1)\cdot(\boldsymbol{u},\boldsymbol{v},t) : (\boldsymbol{u},\boldsymbol{v},t)\in\mathcal{A}\right\}$. Once again, this defines a supporting hyperplane to $\mathcal{A}$, and the dual problem involves finding the supporting hyperplane with *least t*. Since $(\boldsymbol{0},\boldsymbol{0},p^*)\in\mathrm{Bd}(\mathcal{A})$, we have $p^* = (\boldsymbol{0},\boldsymbol{0},p^*)\cdot(\boldsymbol{\lambda},\boldsymbol{\nu},1)\geq g(\boldsymbol{\lambda},\boldsymbol{\nu})$, which is weak duality. *If* there is a non-vertical supporting hyperplane at $(\boldsymbol{0},\boldsymbol{0},p^*)$, then strong duality holds.
- To prove strong duality, let $\mathcal{B} = \left\{(\boldsymbol{0},\boldsymbol{0},s) : s < p^*\right\}$. This set does not intersect with $\mathcal{A}$. Create a separating hyperplane $(\tilde{\boldsymbol{\lambda}},\tilde{\boldsymbol{\nu}},\mu)$ so that
  - $(\tilde{\boldsymbol{\lambda}},\tilde{\boldsymbol{\nu}},\mu)\cdot(\boldsymbol{u},\boldsymbol{v},t)\leq\alpha\ \forall(\boldsymbol{u},\boldsymbol{v},t)\in\mathcal{B}\Rightarrow \mu t\leq\mu p^*\leq\alpha$
  - $(\tilde{\boldsymbol{\lambda}},\tilde{\boldsymbol{\nu}},\mu)\cdot(\boldsymbol{u},\boldsymbol{v},t)\geq\alpha\geq\mu p^*\ \forall(\boldsymbol{u},\boldsymbol{v},t)\in\mathcal{A}\Rightarrow\tilde{\boldsymbol{\lambda}}\geq\boldsymbol{0},\mu\geq 0$. Otherwise, the LHS would be unbounded below as we went up the epigraph.

Divide first equation by μ and substitute in the second to find the Lagrangian is $\geq p^*$.

If $\mu = 0$; second equation above becomes $(\tilde{\boldsymbol{\lambda}}, \tilde{\boldsymbol{\nu}}) \cdot (\boldsymbol{u}, \boldsymbol{v}) \geq 0 \ \forall (\boldsymbol{u}, \boldsymbol{v}, t) \in \mathcal{A}$. Applying this to the strictly feasible point with $\boldsymbol{v} = \boldsymbol{0}$ and $\boldsymbol{u} < \boldsymbol{0}$, we find that $\tilde{\boldsymbol{\lambda}} = 0$, and so $\tilde{\boldsymbol{\nu}} \neq 0$ [supporting hyperplane cannot be $\boldsymbol{0}$ vector], so $\tilde{\boldsymbol{\nu}} \cdot \boldsymbol{v} \geq 0$. But the problem is convex, so $\boldsymbol{v} = A\boldsymbol{x} - \boldsymbol{b}$. But since we have an interior point with $\tilde{\boldsymbol{\nu}} \cdot \boldsymbol{v} = 0$, there are points *around* that interior point with $\tilde{\boldsymbol{\nu}} \cdot \boldsymbol{v} < 0$. Thus, unless $\tilde{\boldsymbol{\nu}}^\top A = \boldsymbol{0}$ [contradiction if A full rank] we have a contradiction. Geometrically, this is equivalent to saying the hyperplane must pass to the left of our interior point. □

- ***Interpretations of Duality***
  - ***Multicriterion optimization***: Consider $\min\left(\boldsymbol{f}(\boldsymbol{x}), f_0(\boldsymbol{x})\right)$. One way to obtain every pareto-optimal point is to minimize $\tilde{\boldsymbol{\lambda}} \cdot \boldsymbol{F}$. Since we can re-scale this without changing the minimizers, $f_0(\boldsymbol{x}) + \boldsymbol{\lambda} \cdot \boldsymbol{F}$ is an example of such a program, which is precisely the Lagrangian.
  - ***Shadow prices***: The dual problem is the lowest cost given we can buy some "constraint violation" and be rewarded when we don't violate them. Clearly, this is lower than our lowest cost without these amenities. When strong duality holds, there is a price that makes us indifferent. This is the "value" of the constraints.
- ***Optimality conditions***
  - ***Complementary Slackness***: Consider that, if $\boldsymbol{x}^*$ and $(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ are primal and dual optimal points with zero duality gap

$$f_0(\boldsymbol{x}^*) = g(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) = \inf_{\boldsymbol{x}} \mathcal{L}(\boldsymbol{x}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) \leq \mathcal{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) \leq f_0(\boldsymbol{x}^*)$$

    As such, every inequality in this line must be an equality. Now, recall that $\mathcal{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) = f_0(\boldsymbol{x}^*) + \boldsymbol{\lambda}^* \cdot \boldsymbol{f}(\boldsymbol{x}^*) + \boldsymbol{\nu}^* \cdot \boldsymbol{g}(\boldsymbol{x}^*)$. This implies that at optimality $\boldsymbol{\lambda}^* \cdot \boldsymbol{f}(\boldsymbol{x}^*) = 0$. Since each term is non-positive, we must have $\lambda_i^* f_i^*(\boldsymbol{x}^*) = 0$. Only active constraints at the optimum can have non-zero multipliers.
  - ***KKT Conditions***: Based on all the above, we find that for *any* problem for which strong duality holds, any primal-dual optimal pair must satisfy

$$\nabla\mathcal{L}(\boldsymbol{x}^*,\boldsymbol{\lambda}^*,\boldsymbol{\nu}^*)=\boldsymbol{0} \qquad \boldsymbol{f}(\boldsymbol{x}^*)\le\boldsymbol{0} \qquad \boldsymbol{h}(\boldsymbol{x}^*)=\boldsymbol{0} \qquad \boldsymbol{\lambda}^*\ge\boldsymbol{0} \qquad \boldsymbol{\lambda}^*\cdot\boldsymbol{f}(\boldsymbol{x}^*)=0$$

If the problem is convex, these are also sufficient conditions, because since $\boldsymbol{\lambda}^*\ge\boldsymbol{0}$, $\mathcal{L}$ is convex, and so the first condition implies $\mathcal{L}$ is minimized. Finally, complementary slackness shows we have 0 duality gap.

Example: $\min\frac{1}{2}\boldsymbol{x}^\top P\boldsymbol{x}+\boldsymbol{q}\cdot\boldsymbol{x}+\boldsymbol{r}$ s.t. $A\boldsymbol{x}=\boldsymbol{b}$. KKT conditions are $A\boldsymbol{x}^*=\boldsymbol{b}$ and $P\boldsymbol{x}^*+\boldsymbol{q}+A^\top\boldsymbol{\nu}^*=\boldsymbol{0}$, or $\begin{bmatrix}P & A^\top\\ A & 0\end{bmatrix}\begin{bmatrix}\boldsymbol{x}^*\\ \boldsymbol{\nu}^*\end{bmatrix}=\begin{bmatrix}-\boldsymbol{q}\\ \boldsymbol{b}\end{bmatrix}$. □

Example: $\min\left(-\sum\log(\alpha_i+x_i)\right)$ s.t. $\boldsymbol{x}\ge 0,\boldsymbol{1}\cdot\boldsymbol{x}=1$. This attempt to maximize communication capacity given total available power of 1 for each channel. KKT:

$$-\tfrac{1}{\alpha_i+x_i}-\lambda_i+\nu=0 \qquad \forall i$$
$$x_i\ge 0,\boldsymbol{1}\cdot\boldsymbol{x}=1 \qquad \lambda_i\ge 0 \qquad \lambda_i x_i=0$$

Clearly, λ_i only acts as a slack variable in the first equation. So

$$\nu\ge\tfrac{1}{\alpha_i+x_i} \qquad \forall i$$
$$x_i\ge 0,\boldsymbol{1}\cdot\boldsymbol{x}=1 \qquad \left(\nu-\tfrac{1}{\alpha_i+x_i}\right)x_i=0$$

The first equation gives $x_i=\frac{1}{\nu}-\alpha_i$, but this can only work if $\alpha_i\le\frac{1}{\nu}\Rightarrow\nu_i\le\frac{1}{\alpha}$. If this were not the case, x_i would go negative. Thus

$$x_i=\begin{cases}\frac{1}{\nu}-\alpha_i & \nu<1/\alpha_i\\ 0 & \nu\ge 1/\alpha_i\end{cases}=\max\left\{0,\tfrac{1}{\nu}-\alpha_i\right\}$$

Using the sum constraint, $\sum_{i=1}^n\max\left\{0,\frac{1}{\nu}-\alpha_i\right\}=1$. This is easy to solve. □

- ***Using the dual to solve the primal***
  - ***Example***: $\min\sum_{i=1}^n f_i(x_i)$ s.t. $\boldsymbol{a}\cdot\boldsymbol{x}=b$, $\mathcal{L}(\boldsymbol{x},\nu)=\sum_{i=1}^n f_i(x_i)+\nu(\boldsymbol{a}\cdot\boldsymbol{x}-b)$. The dual function is $g(\nu)=-\nu b+\inf_{\boldsymbol{x}}\left[\sum_{i=1}^n f_i(x_i)+\nu\boldsymbol{a}\cdot\boldsymbol{x}\right]=-\nu b+\sum_{i=1}^n f_i^*(-\nu a_i)$. The dual therefore involves a single scalar variable (simple to solve). We then use the fact that the optimal point minimizes $\mathcal{L}(\boldsymbol{x},\nu^*)$, which is convex. □
- ***Sensitivity analysis***
  - Consider $\min f_0(\boldsymbol{x})$ s.t. $\boldsymbol{f}(\boldsymbol{x})\le\boldsymbol{u},\boldsymbol{h}(\boldsymbol{x})=\boldsymbol{v}$. We let $p^*(\boldsymbol{u},\boldsymbol{v})$ be its optimal value. If the problem is convex, this is jointly convex (its epigraph is $\mathrm{cl}(\mathcal{A})$, above).
  - ***Global inequality***: Under SD $p^*(\boldsymbol{u},\boldsymbol{v})\ge p^*(\boldsymbol{0},\boldsymbol{0})-\boldsymbol{\lambda}^*\cdot\boldsymbol{u}-\boldsymbol{\nu}^*\cdot\boldsymbol{v}$. To prove,

$$p^*(\boldsymbol{0},\boldsymbol{0}) = g(\boldsymbol{\lambda}^*,\boldsymbol{\nu}^*) \le f_0(\boldsymbol{x}) + \boldsymbol{\lambda}^* \cdot \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{\nu}^* \cdot \boldsymbol{h}(\boldsymbol{x}) \le f_0(\boldsymbol{x}) + \boldsymbol{\lambda}^* \cdot \boldsymbol{u} + \boldsymbol{\nu}^* \cdot \boldsymbol{v}$$

- ***Local Result***: If p differentiable & SD, $\nabla_{\boldsymbol{u}} p^*(\boldsymbol{0},\boldsymbol{0}) = -\boldsymbol{\lambda}^*$ and $\nabla_{\boldsymbol{v}} p^*(\boldsymbol{0},\boldsymbol{0}) = -\boldsymbol{\nu}^*$

  $$\frac{\partial p^*(0,0)}{\partial u_i} = \lim_{t\to 0} \frac{p^*(t\boldsymbol{e}_i,0) - p^*}{t} \ge -\lambda^* \qquad \text{[By global inequality]}$$

  Taking t negative gives the opposite inequality.

- *Examples*
  - Consider $\min \log\left(\sum \exp(\boldsymbol{a}^i \cdot \boldsymbol{x} + b_i)\right)$. The dual isn't particularly interesting. But the dual of $\min \log\left(\sum \exp(y_i)\right)$ s.t. $A\boldsymbol{x} + \boldsymbol{b} = \boldsymbol{y}$ is entropy maximization. The same is true of $\left\| A\boldsymbol{x} - \boldsymbol{b} \right\|$. ☐
  - This can be done with constraints; $f(\boldsymbol{a}\cdot\boldsymbol{x} + b)$ can be transformed to $f(y)$. ☐

- *Theorems of the Alternative*
  - ***Weak alternatives***:
    - ***Non-strict Inequalities***: Consider

      $$\boldsymbol{f}(\boldsymbol{x}) \le \boldsymbol{0}, \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \text{ has sol}^n \Leftrightarrow \min 0 \text{ s.t. } \cdots = 0 \text{ (and not } \infty)$$

      The dual function has the property $g(\alpha\boldsymbol{\lambda}, \alpha\boldsymbol{\nu}) = \alpha g(\boldsymbol{\lambda},\boldsymbol{\nu})$ and is

      $$g(\boldsymbol{\lambda},\boldsymbol{\nu}) = \inf_{\boldsymbol{x}} \left[\boldsymbol{\lambda}\cdot\boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{\nu}\cdot\boldsymbol{h}(\boldsymbol{x})\right]$$

      Because of the homogeneity, if there is *any* $g(\boldsymbol{\lambda},\boldsymbol{\nu}) > 0$ with $\boldsymbol{\lambda} \ge 0$, $d^* = \infty$. If that's infeasible, $d^* = 0$. Thus, since $d^* \le p^*$, we find that

      $$g(\boldsymbol{\lambda},\boldsymbol{\nu}) > 0, \boldsymbol{\lambda} \ge 0 \text{ feasible} \Rightarrow \boldsymbol{f}(\boldsymbol{x}) \le 0, \boldsymbol{h}(\boldsymbol{x}) = 0 \text{ infeasible}$$

      In fact, *at most* one of the two is feasible – *weak alternatives*.
    - ***Strict inequalities***: If the inequality $\boldsymbol{f}(\boldsymbol{x}) < \boldsymbol{0}$ is strict, the alternative is

      $$g(\boldsymbol{\lambda},\boldsymbol{\nu}) \ge 0, \boldsymbol{\lambda} > 0 \text{ feasible} \Rightarrow \boldsymbol{f}(\boldsymbol{x}) < 0, \boldsymbol{h}(\boldsymbol{x}) = 0 \text{ infeasible}$$

      We can show this directly from the definition of the dual function, if we assume there exists $\boldsymbol{f}(\tilde{\boldsymbol{x}}) < \boldsymbol{0}$, then $\exists \tilde{\boldsymbol{\lambda}} > 0, \boldsymbol{\nu}$ s.t. $g(\boldsymbol{\lambda},\boldsymbol{\nu}) < 0$.
  - ***Strong alternatives*** – when f_i are convex and h_i are affine, we might be able to prove *strong alternatives*; that *exactly* one of them must hold
    - ***Strict inequalities***: First, consider

      $$\boldsymbol{f}(\boldsymbol{x}) < \boldsymbol{0}, A\boldsymbol{x} = \boldsymbol{b} \text{ feasible} \Leftrightarrow p^* = \left(\min s \text{ s.t. } f_i(\boldsymbol{x}) \le s, A\boldsymbol{x} = \boldsymbol{b}\right) < 0$$

The dual function is $g(\boldsymbol{\lambda},\boldsymbol{\nu}) = \boldsymbol{\nu}\cdot(A\boldsymbol{x}-\boldsymbol{b}) + \min_s\left[(1-\boldsymbol{1}\cdot\boldsymbol{\lambda})s\right]$. This is only finite if $\boldsymbol{1}\cdot\boldsymbol{\lambda}=1$. So the dual is $d^* = \left(\max g(\boldsymbol{\lambda},\boldsymbol{\nu}) \text{ s.t. } \boldsymbol{1}\cdot\boldsymbol{\lambda}=1, \boldsymbol{\lambda}\geq\boldsymbol{0}\right)$. Provided strict feasibility holds, strong duality holds and $p^*=d^*$. So if the original system is infeasible ($p^* \geq 0$), then there exists a $g(\boldsymbol{\lambda},\boldsymbol{\nu})\geq 0, \boldsymbol{\lambda}>\boldsymbol{0}$. Similarly, if there exists such a $(\boldsymbol{\lambda},\boldsymbol{\nu})$, then $p^*\geq 0$…

$$\boldsymbol{f}(\boldsymbol{x})<\boldsymbol{0}, A\boldsymbol{x}=\boldsymbol{b} \ \frac{\text{feasible}}{\text{infeasible}} \Leftrightarrow g(\boldsymbol{\lambda},\boldsymbol{\nu})\geq 0, \boldsymbol{\lambda}>\boldsymbol{0} \ \frac{\text{infeasible}}{\text{feasible}}$$

- ***Non-strict inequalities***: Consider $\boldsymbol{f}(\boldsymbol{x})\leq\boldsymbol{0}, A\boldsymbol{x}=\boldsymbol{b}$ the program is the same as above, but we need the optimum to be attained so that $p^*>0$ if the system is infeasible. In that case, $\boldsymbol{\lambda}\geq\boldsymbol{0}, g(\boldsymbol{\lambda},\boldsymbol{\nu})>0$ is clearly feasible.

- ***Example***: Consider $A\boldsymbol{x}\leq\boldsymbol{b}$. Then $g(\boldsymbol{\lambda}) = -\boldsymbol{\lambda}\cdot\boldsymbol{b}$ if $A^\top\boldsymbol{\lambda}=\boldsymbol{0}$ and $-\infty$ o.w. The strong system of alternative inequalities is $\boldsymbol{\lambda}\geq\boldsymbol{0}, A^\top\boldsymbol{\lambda}=0, \boldsymbol{\lambda}\cdot\boldsymbol{b}<0$. □
- ***Example***: Take m ellipsoids $\mathcal{E}_i = \left\{\boldsymbol{x} : f_i(\boldsymbol{x}) = \boldsymbol{x}^\top A_i\boldsymbol{x} + 2\boldsymbol{b}^i\cdot\boldsymbol{x} + \boldsymbol{c}^i \leq 0\right\}, A_i\in\mathbb{S}^n_{++}$. We ask if the intersection has a non-empty interior. This is equivalent to solving the system $\boldsymbol{f}(\boldsymbol{x})<\boldsymbol{0}$. Here, $g(\boldsymbol{\lambda}) = \inf_{\boldsymbol{x}} \boldsymbol{x}^\top\left(\sum\lambda_i A_i\right)\boldsymbol{x} + 2\left(\sum\lambda_i\boldsymbol{b}^i\right)\cdot\boldsymbol{x} + \left(\sum\lambda_i\boldsymbol{c}^i\right)$. Differentiating, setting to 0 and using obvious notation, $g(\boldsymbol{\lambda}) = -\boldsymbol{b}_\lambda^\top A_\lambda^{-1}\boldsymbol{b}_\lambda + \boldsymbol{c}_\lambda$. As such, the alternative system is $\boldsymbol{\lambda}>\boldsymbol{0}, -\boldsymbol{b}_\lambda^\top A_\lambda^{-1}\boldsymbol{b}_\lambda + \boldsymbol{c}_\lambda \geq \boldsymbol{0}$.

  To explain geometrically, consider that the ellipsoid with $f(\boldsymbol{x}) = \boldsymbol{\lambda}\cdot\boldsymbol{f}(\boldsymbol{x})$ contains the intersection of all the ellipsoids above, because if $\boldsymbol{f}(\boldsymbol{x})\leq\boldsymbol{0}$, then clearly a positive linear combination of them is also ≤ 0. This ellipsoid is empty if and only if the alternative is satisfied [prove by finding $\inf f(\boldsymbol{x})$]. □
- ***Example***: Farkas' Lemma: the following two systems are strong alternatives

  $A\boldsymbol{x}=\boldsymbol{b}, \boldsymbol{x}\geq\boldsymbol{0}$ $\qquad\qquad$ $A^\top\boldsymbol{y}\geq 0, \boldsymbol{y}\cdot\boldsymbol{b}<0$
- ………………..

- ***Duality & Decentralization***
  - Consider $\min\sum_{i=1}^k f_i(\boldsymbol{x}^i)$ s.t. $\sum_{i=1}^k \boldsymbol{g}^i(\boldsymbol{x}^i)\leq\boldsymbol{0}, \boldsymbol{x}^i\in\Omega_i$ [note: the vector $\boldsymbol{g}$ represents a number of inequality constraints]. The Lagrangian is $\mathcal{L}(\boldsymbol{x},\boldsymbol{\mu}) = \sum_{i=1}^k f_i(\boldsymbol{x}^i) + \boldsymbol{\mu}\cdot\sum_{i=1}^k \boldsymbol{g}^i(\boldsymbol{x}^i)$. The dual is $g(\boldsymbol{\mu}) = \sum_{i=1}^k g_i(\boldsymbol{\mu})$ s.t. $\boldsymbol{\mu}\geq\boldsymbol{0}$ where $g_i(\boldsymbol{\mu}) = \inf_{\boldsymbol{x}^i\in\Omega_i} f_i(\boldsymbol{x}^i) + \boldsymbol{\mu}\cdot\boldsymbol{g}^i(\boldsymbol{x}^i)$

- For example x_{ij} could be the quantity of resource j allocated to acitivity i, and we might want to *maximize* utility. Each component of $\boldsymbol{g}^i$ corresponds to one resource constraint. The resulting geometric multipliers can be considered as *prices* for a given resources.
- The Tatonnement procedure guesses initial prices, solves the problem, and then adjusts them to ensure each resource is used exactly.

- ***Duality & Combinatorial Optimization***
  - The knapsack problem is $\max\ \boldsymbol{v}\cdot\boldsymbol{x}$ s.t. $\boldsymbol{w}\cdot\boldsymbol{x}\le C, \boldsymbol{x}\in\{0,1\}^n$.
    - $g(\mu)=\max_{\boldsymbol{x}\in\{0,1\}^n}(\boldsymbol{v}-\mu\boldsymbol{w})\cdot\boldsymbol{x}+\mu C=\mu C+\sum w_i\max([v_i/w_i]-\mu,0)$
    - Assume WLOG $\frac{v_1}{w_1}\ge\cdots\ge\frac{v_n}{w_n}$. There is then a breakpoint $I^*(\mu)$ until which $v_i/w_i\ge\mu$. We can write $g(\mu)=\mu C+\sum_{i=1}^{I^*(\mu)}v_i-\mu w_i$
    - The dual problem is to make this as small as possible for $\mu\ge 0$. Consider that $g(\mu)=\left(\sum_{i=1}^{I^*(\mu)}v_i\right)+\mu\left(C-\sum_{i=1}^{I^*(\mu)}w_i\right)$. This is piecewise-linear in μ – the minimum occurs when the gradient switches from negative to positive; so $I^{\text{opt}}=\min\left\{I:\sum_{i=1}^{I}w_i>C\right\}$. We then want μ to take its smallest possible value, which isd $\mu^*=v_{I^{\text{opt}}}/w_{I^{\text{opt}}}$
    - This gives an upper bound $\left(\sum_{i=1}^{I^{\text{opt}}}v_i\right)+\mu^*\left(C-\sum_{i=1}^{I^{\text{opt}}}w_i\right)$.
  - For a lower-bound, consider the solution with $\tilde{x}_i=1$ if $i<I^{\text{opt}}$ and 0 otherwise. This is clearly feasible, and corresponds to the greatest "bang for buck" policy.

A More Rigorous Approach to Optimality Conditions

- ***Unconstrained optimization***: interior solutions of $\min f(\boldsymbol{x})$ s.t. $\boldsymbol{x}\in\mathcal{C}\subset\mathbb{R}^n$
  - ***Necessary conditions***: If $\boldsymbol{x}^*$ is in the interior of the feasible region and is an optimal minimum, then $\nabla f(\boldsymbol{x}^*)=0$ and $\nabla^2 f(\boldsymbol{x}^*)\succeq 0$.
  - ***Sufficient conditions***: If $\boldsymbol{x}^*$ is in the interior of the feasible region and $\nabla f(\boldsymbol{x}^*)=0$ and $\nabla^2 f(\boldsymbol{x}^*)\succ 0$, then $\boldsymbol{x}^*$ is a strict local minimum.

- When using these conditions, (1) verify existence (2) find points with $\nabla f(\boldsymbol{x}) = \boldsymbol{0}$ (3) compare those to points on the boundary.
- Consider, instead $f(\boldsymbol{x}, \boldsymbol{a})$. Differentiate the FOCs with respect to $\boldsymbol{a}$ to get

$$\nabla \boldsymbol{x}^*(\boldsymbol{a}) = -\nabla^2_{xa} f(\boldsymbol{x}^*(\boldsymbol{a}), \boldsymbol{a}) \left\{ \nabla^2_{xx} f(\boldsymbol{x}^*(\boldsymbol{a}), \boldsymbol{a}) \right\}^{-1}$$

$$\nabla f^*(\boldsymbol{a}) = \nabla_a f(\boldsymbol{x}^*(\boldsymbol{a}), \boldsymbol{a}) + \nabla \boldsymbol{x}^*(\boldsymbol{a}) \nabla_x f(\boldsymbol{x}^*(\boldsymbol{a}), \boldsymbol{a})$$

- ***Constrained optimization***: boundary solutions of $\min f(\boldsymbol{x})$ s.t. $\boldsymbol{x} \in \mathcal{C} \subset \mathbb{R}^n$.
  - The *set of descent directions* is $\mathcal{D}(\boldsymbol{x}^*) = \left\{ \boldsymbol{d} \in \mathbb{R}^n : \nabla f(\boldsymbol{x}^*) \cdot \boldsymbol{d} < 0 \right\}$. The *tangent cone* $\mathcal{T}(\boldsymbol{x}^*)$ is the set of directions we can move from $\boldsymbol{x}^*$ while staying in the feasible region. A necessary condition for $\boldsymbol{x}^*$ to be optimal is $\mathcal{D}(\boldsymbol{x}^*) \cap \mathcal{T}(\boldsymbol{x}^*) = \varnothing$.
  - If $\mathcal{C}$ is defined only by the equality constraints $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}$, then for any *regular point* (ie: point at which the gradients $\nabla h_i(\boldsymbol{x}^*)$ are linearly independent):

$$\mathcal{T}(\boldsymbol{x}^*) = \mathcal{V}(\boldsymbol{x}^*) = \left\{ \boldsymbol{d} \in \mathbb{R}^n : \nabla \boldsymbol{h}(\boldsymbol{x}^*)^\top \boldsymbol{d} = \boldsymbol{0} \right\} = \mathcal{N}\left(\nabla \boldsymbol{h}(\boldsymbol{x}^*)^\top \right)$$

    Intuitively, any move in direction $\boldsymbol{d}$ will change $\boldsymbol{h}$ by $\nabla \boldsymbol{h}(\boldsymbol{x}^*)^\top \boldsymbol{d}$… Thus

$$\boldsymbol{x}^* \text{ local minimum \& regular} \Rightarrow \nabla f(\boldsymbol{x}^*) \in \mathcal{V}(\boldsymbol{x}^*)^\perp = \mathcal{N}\left(\nabla \boldsymbol{h}(\boldsymbol{x}^*)^\top \right)^\top = \mathcal{R}\left(\nabla \boldsymbol{h}(\boldsymbol{x}^*) \right)$$

    [For $\boldsymbol{d} \in \mathcal{V}$, we need $\nabla f(\boldsymbol{x}^*) \cdot \boldsymbol{d} \geq 0$. But since $\boldsymbol{d} \in \mathcal{V} \Rightarrow -\boldsymbol{d} \in \mathcal{V}$, $\nabla f(\boldsymbol{x}^*) \cdot \boldsymbol{d} = 0$].

    Thus, for any regular local optimum $\boldsymbol{x}^*$, $\exists_{\text{unique}}$ $\boldsymbol{\lambda}^*$ s.t. $\nabla f(\boldsymbol{x}^*) + \boldsymbol{\lambda} \cdot \nabla \boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}$.

    Example: $\min\ \boldsymbol{x}^\top \Gamma \boldsymbol{x}$ s.t. $\boldsymbol{1} \cdot \boldsymbol{x} = 1, \boldsymbol{\mu} \cdot \boldsymbol{x} = \bar{\mu}$. FOCs are

$$2\Gamma \boldsymbol{x} + \lambda_1 \boldsymbol{1} + \lambda_2 \boldsymbol{\mu} = 0 \qquad \boldsymbol{1} \cdot \boldsymbol{x}^* = 1 \qquad \boldsymbol{\mu}^\top \boldsymbol{x}^* = \bar{\mu}$$

    The first gives $\boldsymbol{x}^* = -\frac{1}{2}\Gamma^{-1}\left(\lambda_1 \boldsymbol{1} + \lambda_2 \boldsymbol{\mu} \right)$. Feeding into the others gives a system of equations for $\boldsymbol{\lambda}$, whence $\boldsymbol{\lambda} = \boldsymbol{\eta} + \boldsymbol{\zeta}\bar{\mu} \Rightarrow \boldsymbol{x}^* = \bar{\mu}\boldsymbol{v} + \boldsymbol{w} \Rightarrow \sigma^2 = (\alpha\bar{\mu} + \beta)^2 + \gamma$. ☐
  - Consider the addition of inequality constraints $\boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}$ to the definition of $\mathcal{C}$. It can be shown that all constraints but the active ones at the optimum can be ignored. Thus, provided a point is regular (ie: $\left\{ \nabla h_i(\boldsymbol{x}^*) \right\} \cup \left\{ \nabla g_j(\boldsymbol{x}^*) : j \text{ active} \right\}$ is linearly independent), the KKT conditions provide conditions for optimality.
  - When using such conditions, it is important to check for non-regular points as well. Constraint qualifications can be weakened to requiring inequalities to be convex and equalities to be linear.
- ***Subgradients*** – another way of expressing optimality conditions is as follows

$$\begin{aligned} \boldsymbol{x} \in \arg\min\left\{f(\boldsymbol{x}) : \boldsymbol{x} \in \mathcal{C}\right\} &\Leftrightarrow \boldsymbol{x} \in \arg\min\left\{f(\boldsymbol{x}) + \mathbb{1}_{\mathcal{C}}(\boldsymbol{x})\right\} \\ &\Leftrightarrow \boldsymbol{0} \in \partial\left[f(\boldsymbol{x}) + \mathbb{1}_{\mathcal{C}}(\boldsymbol{x})\right] = \partial f(\boldsymbol{x}) + \partial \mathbb{1}_{\mathcal{C}}(\boldsymbol{x}) \\ &\Leftrightarrow \exists \boldsymbol{g} \in \partial f(\boldsymbol{x}) \text{ s.t. } -\boldsymbol{g} \in \partial \mathbb{1}_{\mathcal{C}}(\boldsymbol{x}) \\ &\Leftrightarrow \exists \boldsymbol{g} \in \partial f(\boldsymbol{x}) \text{ s.t. } \boldsymbol{g} \cdot (\boldsymbol{y} - \boldsymbol{x}) \geq 0 \ \forall \boldsymbol{y} \in \mathcal{C} \end{aligned}$$

For the last step, note $\boldsymbol{g} \in \partial \mathbb{1}_{\mathcal{C}}(\boldsymbol{x}) \Leftrightarrow \mathbb{1}_{\mathcal{C}}(\boldsymbol{y}) \geq \mathbb{1}_{\mathcal{C}}(\boldsymbol{x}) + \boldsymbol{g} \cdot (\boldsymbol{y} - \boldsymbol{x}) \Leftrightarrow 0 \geq \boldsymbol{g} \cdot (\boldsymbol{y} - \boldsymbol{x}) \ \forall \boldsymbol{y} \in \mathcal{C}$.

Chapter 6 – Approximation

- The most basic approximation problem is $\min \left\| A\boldsymbol{x} - \boldsymbol{b} \right\|$. Has solution 0 if $\boldsymbol{b} \in \mathcal{R}(A)$.
  - Approximating $\boldsymbol{b}$ as closely as possible using columns of A.
  - Letting $\boldsymbol{y} = A\boldsymbol{x} + \boldsymbol{v}$ ($\boldsymbol{v}$ is noise), and guessing $\boldsymbol{x}$ based on $\boldsymbol{y}$, assuming noise small
  - $\min_{\boldsymbol{u} \in \mathcal{R}(A)} \left\| \boldsymbol{u} - \boldsymbol{b} \right\|$; projecting $\boldsymbol{b}$ onto $\mathcal{R}(A)$.
  - $\boldsymbol{x}$ are design variables, $\boldsymbol{b}$ is a target, $A\boldsymbol{x}$ is the result.
- ***Examples***
  - $\min \left\| A\boldsymbol{x} - \boldsymbol{b} \right\|_2^2$; least square. Solution $A^\top A\boldsymbol{x} = A^\top \boldsymbol{b}$.
  - $\min \left\| A\boldsymbol{x} - \boldsymbol{b} \right\|_\infty^2$; *Chebyshev approx problem*. Same as $\min t$ s.t. $-t\boldsymbol{1} \leq A\boldsymbol{x} - \boldsymbol{b} \leq t\boldsymbol{1}$
  - $\min \left\| A\boldsymbol{x} - \boldsymbol{b} \right\|_1^2$; *Robust estimator*. Same as $\min \boldsymbol{1} \cdot \boldsymbol{t}$ s.t. $-\boldsymbol{t} \leq A\boldsymbol{x} - \boldsymbol{b} \leq \boldsymbol{t}$. Slowest growing that is still convex.
  - $\min \phi(r_1) + \cdots + \phi(r_n)$ s.t. $\boldsymbol{r} = A\boldsymbol{x} - \boldsymbol{b}$ is *penalty function approximation*. Measure of dislike of large residuals. ☐
- ***Least norm problems*** are $\min \left\| \boldsymbol{x} \right\|$ s.t. $A\boldsymbol{x} = \boldsymbol{b}$. Can be reformulated as $\min \left\| \boldsymbol{x}_0 + Z\boldsymbol{u} \right\|$, where $\boldsymbol{x}_0$ is a solution and the columns of Z form a basis for $\mathcal{N}(A)$.
- ***Regularization problems*** are bi-objective problems; $\min \left(\left\| A\boldsymbol{x} - \boldsymbol{b} \right\|, \left\| \boldsymbol{x} \right\| \right)$. Use $\min \left\| A\boldsymbol{x} - \boldsymbol{b} \right\| + \gamma \left\| \boldsymbol{x} \right\|$ to trace out the tradeoff curve.
  - ***Example (Tikhonov regularization)***:

$$\min \left\| A\boldsymbol{x} - \boldsymbol{b} \right\|_2^2 + \gamma \left\| \boldsymbol{x} \right\|_2^2 = \boldsymbol{x}^\top \left(A^\top + \gamma I \right) \boldsymbol{x} - 2\boldsymbol{b}^\top A\boldsymbol{x} + \boldsymbol{b}^\top \boldsymbol{b} = \left\| \begin{pmatrix} A \\ I\sqrt{\gamma} \end{pmatrix} \boldsymbol{x} - \begin{pmatrix} \boldsymbol{b} \\ \boldsymbol{0} \end{pmatrix} \right\|_2$$

Solution is

$$\begin{pmatrix} A \\ I\sqrt{\gamma} \end{pmatrix}^\top \begin{pmatrix} A \\ I\sqrt{\gamma} \end{pmatrix} \boldsymbol{x} = \begin{pmatrix} A \\ I\sqrt{\gamma} \end{pmatrix}^\top \begin{pmatrix} \boldsymbol{b} \\ \boldsymbol{0} \end{pmatrix} \Rightarrow \boldsymbol{x} = \left(A^\top A + \gamma I\right)^{-1} A^\top \boldsymbol{b} \qquad \square$$

- o ***Example (minimizing derivatives and curvature)***: We can replace our second objective ($\|\boldsymbol{x}\|$) by $\|D\boldsymbol{x}\|$. We then get $\boldsymbol{x}^* = (A^\top A + \gamma D^\top D)^{-1} A^\top \boldsymbol{b}$. Two useful examples of D are
  - D has 1s across its diagonal and a –1 to the left of each 1. $D\boldsymbol{x}$ is then the vector of quantities $x_{i+1} - x_i$; the "discrete derivative" of $\boldsymbol{x}$.
  - D has 2s across its diagonal, and 1s to its left and right. $D\boldsymbol{x}$ is then the vector of quantities $\left([x_{i+1} - x_i] - [x_i - x_{i-1}]\right) = x_{i+1} - 2x_i + x_{i-1}$, approximately the *curvature* (second derivative) of $\boldsymbol{x}$. $\square$
- o ***Example (LASSO)***: $\min \|A\boldsymbol{x} - \boldsymbol{b}\|_2^2 + \gamma \|\boldsymbol{x}\|_1$ can be written as QP

$$\min \|A\boldsymbol{x} - \boldsymbol{b}\|_2^2 + \gamma \boldsymbol{1} \cdot \boldsymbol{y} \text{ s.t. } -\boldsymbol{y} \le \boldsymbol{x} \le \boldsymbol{y} \qquad \square$$

- ***Stochastic Robust approximation***: $\min_{\boldsymbol{x}} \|A\boldsymbol{x} - \boldsymbol{b}\|$ with an uncertain A that has a probability distribution. Do $\min_{\boldsymbol{x}} \mathbb{E}\|A\boldsymbol{x} - \boldsymbol{b}\|$ instead. If $\mathbb{P}(A = A_i) = p_i$, this becomes $\min \boldsymbol{p} \cdot \boldsymbol{t}$ s.t. $\|A_i\boldsymbol{x} - \boldsymbol{b}\| \le t_i$. For a 2-norm, this is an SOCP. For a 1-norm or ∞-norm, this can be written as an LP.

  $\min \mathbb{E}\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2$ is actually tractable. We can write $A = \bar{A} + U$, and we can then write the objective as $\min\left(\|\bar{A}\boldsymbol{x} - \boldsymbol{b}\|_2^2 + \|P^{1/2}\boldsymbol{x}\|_2^2\right)$, with $P = \mathbb{E}(U^\top U)$. This is Tikhonov regularized least-squares.
- ***Worst-case Robust approximation***: We let $A \in \mathcal{A}$, and we solve $\min_{\boldsymbol{x}} \sup_{A\in\mathcal{A}} \|A\boldsymbol{x} - \boldsymbol{b}\|$. We consider several sets $\mathcal{A}$:
  - o ***Finite set or polyhedron***: $\mathcal{A} = \{A_1, \cdots, A_k\}$: $\min t$ s.t. $\|A_i\boldsymbol{x} - \boldsymbol{b}\| \le t \ \forall i$. For a polyhedron, try this out for the vertices.
  - o ***Norm bound error***: $\mathcal{A} = \{\bar{A} + U \mid \|U\| \le a\}$, where the norm is a matrix-norm. Consider the approximation problem with the Euclidean norm and the max-singular-value norm. Then $\sup_{A\in\mathcal{A}} \|\bar{A}\boldsymbol{x} - \boldsymbol{b} + U\boldsymbol{x}\|$ occurs when $U\boldsymbol{x}$ is aligned with $\bar{A}\boldsymbol{x} - \boldsymbol{b}$ and is as large as can be. Letting $U = a(\bar{A}\boldsymbol{x} - \boldsymbol{b})\boldsymbol{x}^\top / \|\bar{A}\boldsymbol{x} - \boldsymbol{b}\|_2 \|\boldsymbol{x}\|_2$

achieves that. Thus, our program is $\min_{\boldsymbol{x}} \left\|\bar{A}\boldsymbol{x} - \boldsymbol{b}\right\|_2 + a\left\|\boldsymbol{x}\right\|_2$. This is solvable as an SOCP: $\min t_1 + at_2$ s.t. $\left\|\bar{A}\boldsymbol{x} - \boldsymbol{b}\right\|_2 \le t_1, \left\|\boldsymbol{x}\right\|_2 \le t_2$.

- ***Uncertainty ellipsoids***: Assume each row of A is in $\mathcal{E}_i = \{\bar{\boldsymbol{A}}_{\text{row } i} + P_i\boldsymbol{u} \mid \left\|\boldsymbol{u}\right\|_2 \le 1\}$. Then $\sup_{A\in\mathcal{A}} \left\|A\boldsymbol{x} - \boldsymbol{b}\right\|$ can be found by individually maximizing $\sup\left|\boldsymbol{A}_{\text{row } i} \cdot \boldsymbol{x} - b_i\right| = \sup_{\boldsymbol{u}} \left\{\left|\bar{\boldsymbol{A}}_{\text{row } i} \cdot \boldsymbol{x} - b_i + (P_i\boldsymbol{u})\cdot\boldsymbol{x}\right| \mid \left\|\boldsymbol{u}\right\|_2 \le 1\right\}$. We do this by aligning $\boldsymbol{u}$ with $P_i^\top \boldsymbol{x}$ and get $\sup\left|\boldsymbol{A}_{\text{row } i} \cdot \boldsymbol{x} - b_i\right| = \left|\bar{\boldsymbol{A}}_{\text{row } i} \cdot \boldsymbol{x} - b_i\right| + \left\|P_i^\top \boldsymbol{x}\right\|_2$. We then have $\min_{\boldsymbol{x}} \sup_{A\in\mathcal{A}} \left\|A\boldsymbol{x} - \boldsymbol{b}\right\| = \min_{\boldsymbol{x},\boldsymbol{t}} \left\|\boldsymbol{t}\right\|_2$ s.t. $\left|\bar{\boldsymbol{A}}_{\text{row } i} \cdot \boldsymbol{x} - b_i\right| + \left\|P_i^\top \boldsymbol{x}\right\|_2 \le t_i$. We can get rid of the absolute value sign and put the problem in epigraph form to get an SOCP.

Chapter 7 – Statistical Estimation

- In a parametric estimation problem, we believe that a random variable of interest has a density that is part of a family $p_{\boldsymbol{x}}(\boldsymbol{y})$ indexed by $\boldsymbol{x} \in \Omega$. Given an observation $\boldsymbol{y}$, the *maximum-likelihood* estimation problem is $\max \log p_{\boldsymbol{x}}(\boldsymbol{y})$ s.t. $\boldsymbol{x} \in \Omega$.
- ***Example***: Consider a model $y_i = \boldsymbol{a}^i \cdot \boldsymbol{x} + v_i$, where y_i are the observed quantities, $\boldsymbol{x}$ is to be estimated and the v_i are IID noises with a density p. Then $p_{\boldsymbol{x}}(\boldsymbol{y}) = \prod_{i=1}^m p(y_i - \boldsymbol{a}^i \cdot \boldsymbol{x})$
  - ***Gaussian noise***: $\ell(\boldsymbol{x}) = -\frac{m}{2}\log(2\pi\sigma^2) - \frac{1}{2\sigma^2}\left\|A\boldsymbol{x} - \boldsymbol{y}\right\|_2^2$. ML is least-squares.
  - ***Laplacian noise***: $p(\boldsymbol{z}) = e^{-|\boldsymbol{z}|/\alpha} / 2\alpha$, and $\ell(\boldsymbol{x}) = -m\log 2\alpha - \left\|\boldsymbol{a}^i \cdot \boldsymbol{x} - y_i\right\|_1 / \alpha$. This is $\ell_1 -$norm approximation.
  - ***Uniform*** $[-\boldsymbol{\alpha}, \boldsymbol{\alpha}]$ ***noise***: $\ell(\boldsymbol{x}) = -m\log 2\alpha$ if $\left|y_i - \boldsymbol{a}^i \cdot \boldsymbol{x}\right| \le \alpha\ \forall i$ and $-\infty$ otherwise. The ML estimate is then any $\boldsymbol{x}$ with $\left|y_i - \boldsymbol{a}^i \cdot \boldsymbol{x}\right| \le \alpha\ \forall i$.
- ***Example (logistic regression)***: We let y_i be Bernoulli random variables with $p_i = \exp(\boldsymbol{a}\cdot\boldsymbol{u}^i + \boldsymbol{b}) / [1 + \exp(\boldsymbol{a}\cdot\boldsymbol{u}^i + \boldsymbol{b})]$. Then $\ell_{\boldsymbol{a},\boldsymbol{b}} = \sum_{y_i=1} \log p_i + \sum_{y_i=0} \log(1 - p_i)$. We can feed p_i into this expression.

Chapter 8 – Geometric problems

- In a *classification problem*, we are given two sets of points $\left\{\boldsymbol{x}^1,\cdots,\boldsymbol{x}^N\right\}$ and $\left\{\boldsymbol{y}^1,\cdots,\boldsymbol{y}^M\right\}$ and wish to find a function f (within a family) s.t. $f(\boldsymbol{x}^i)>0\ \forall i$ and $f(\boldsymbol{y}^i)<0\ \forall i$.
- ***Linear discrimination***: We look for a function $f(\boldsymbol{x})=\boldsymbol{a}\cdot\boldsymbol{x}-b$ such that $f(\boldsymbol{x}^i)\geq 1$ and $f(\boldsymbol{y}^i)\leq -1$ (where we have simply normalized the equations above) [ie: we seek a hyperplane that separates the two sets of points].

  Interestingly, the strong alternative of this system of equations is
  $$\boldsymbol{\lambda}\geq\boldsymbol{0},\tilde{\boldsymbol{\lambda}}\geq\boldsymbol{0},(\boldsymbol{\lambda},\tilde{\boldsymbol{\lambda}})\neq\boldsymbol{0},\sum\lambda_i\boldsymbol{x}^i=\sum\tilde{\lambda}_i\boldsymbol{y}^i,\boldsymbol{1}\cdot\boldsymbol{\lambda}=\boldsymbol{1}\cdot\tilde{\boldsymbol{\lambda}}$$
  By dividing by $\boldsymbol{1}\cdot\boldsymbol{\lambda}$, this becomes $\boldsymbol{\lambda},\tilde{\boldsymbol{\lambda}}\geq\boldsymbol{0},\boldsymbol{1}\cdot\boldsymbol{\lambda}=1,\boldsymbol{1}\cdot\tilde{\boldsymbol{\lambda}}=1,\sum\lambda_i\boldsymbol{x}^i=\sum\tilde{\lambda}_i\boldsymbol{y}^i$. This states that the convex hulls of $\boldsymbol{x}$ and $\boldsymbol{y}$ intersect.
- ***Robust linear discrimination***: There are lots of possible solutions to the problem above; we'd like to find the one that separates the points *the most*. Consider the planes $\left\{\boldsymbol{a}\cdot\boldsymbol{z}+b=1\right\}$ and $\left\{\boldsymbol{a}\cdot\boldsymbol{z}+b=-1\right\}$. To find the distance between them, take a point with $\boldsymbol{a}\cdot\boldsymbol{x}+b=1$ and solve $\boldsymbol{a}\cdot(\boldsymbol{x}+t\boldsymbol{a})+b=-1\Rightarrow t\left\|\boldsymbol{a}\right\|_2^2=-2\Rightarrow$ Distance $=2/\left\|\boldsymbol{a}\right\|_2^2$. So, we want to solve $\min\frac{1}{2}\left\|\boldsymbol{a}\right\|_2$ s.t. $f(\boldsymbol{x}^i)\geq t\ \forall i, f(\boldsymbol{y}^i)\leq -t\ \forall i$.

  The dual function is $g(\boldsymbol{\lambda},\boldsymbol{\nu})=\frac{1}{2}\left\|\boldsymbol{a}\right\|_2+t(\boldsymbol{\lambda}\cdot\boldsymbol{1})+t(\boldsymbol{\nu}\cdot\boldsymbol{1})+\left[\boldsymbol{\nu}\cdot\boldsymbol{1}-\boldsymbol{\lambda}\cdot\boldsymbol{1}\right]b+(\boldsymbol{\nu}^\top\boldsymbol{Y}-\boldsymbol{\lambda}^\top\boldsymbol{X})\boldsymbol{a}$. We need $\boldsymbol{1}\cdot\boldsymbol{\nu}=\boldsymbol{1}\cdot\boldsymbol{\lambda}$, in which case $g(\boldsymbol{\lambda},\boldsymbol{\nu})=\frac{1}{2}\left\|\boldsymbol{a}\right\|_2+(\boldsymbol{\nu}^\top\boldsymbol{Y}-\boldsymbol{\lambda}^\top\boldsymbol{X})\boldsymbol{a}+t(\boldsymbol{\lambda}\cdot\boldsymbol{1}+\boldsymbol{\nu}\cdot\boldsymbol{1})$. We then note that by Cauchy-Schwarz, $\left|(\boldsymbol{\nu}^\top\boldsymbol{Y}-\boldsymbol{\lambda}^\top\boldsymbol{X})\boldsymbol{a}\right|\leq\left\|\boldsymbol{\nu}^\top\boldsymbol{Y}-\boldsymbol{\lambda}^\top\boldsymbol{X}\right\|_2\left\|\boldsymbol{a}\right\|_2$ and so the dual is $\max\boldsymbol{\lambda}\cdot\boldsymbol{1}+\boldsymbol{\nu}\cdot\boldsymbol{1}$ s.t. $\boldsymbol{1}\cdot\boldsymbol{\nu}=\boldsymbol{1}\cdot\boldsymbol{\lambda},\left\|\boldsymbol{\nu}^\top\boldsymbol{Y}-\boldsymbol{\lambda}^\top\boldsymbol{X}\right\|\leq\frac{1}{2},\boldsymbol{\lambda},\boldsymbol{\nu}\geq\boldsymbol{0}$. Normalizing $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ so that $\boldsymbol{1}\cdot\boldsymbol{\lambda}=\boldsymbol{1}\cdot\boldsymbol{\nu}=1$, we get $\min t$ s.t. $\left\|\boldsymbol{\nu}^\top\boldsymbol{Y}-\boldsymbol{\lambda}^\top\boldsymbol{X}\right\|\leq t,\boldsymbol{\lambda},\boldsymbol{\mu}\geq\boldsymbol{0},\boldsymbol{1}\cdot\boldsymbol{\lambda}=\boldsymbol{1}\cdot\boldsymbol{\nu}=1$. This is the minimum distance between the convex hull of the points.

  Practically, we would minimize the program above using $\left\|\boldsymbol{x}\right\|_2^2$. The dual is relatively simple to construct as a QP, and the primal solution can be recovered: $\boldsymbol{a}=\boldsymbol{\lambda}^\top\boldsymbol{X}-\boldsymbol{\nu}^\top\boldsymbol{Y}$
- ***Approximate linear separation***: If the points cannot be exactly separated, we might try to solve $\min\boldsymbol{1}\cdot\boldsymbol{u}+\boldsymbol{1}\cdot\boldsymbol{v}$ s.t. $\boldsymbol{a}\cdot\boldsymbol{x}^i-b\geq 1-u_i\ \forall i,\boldsymbol{a}\cdot\boldsymbol{y}^i-b\leq -(1-v_i)\ \forall i,\boldsymbol{u}\geq\boldsymbol{0},\boldsymbol{v}\geq\boldsymbol{0}$. This is a heuristic to minimize the number of misclassifications.

  The *support vector classifier* minimizes a *trade-off* between $\boldsymbol{1}\cdot\boldsymbol{u}+\boldsymbol{1}\cdot\boldsymbol{v}$ and $\frac{1}{2}\left\|\boldsymbol{a}\right\|_2^2$. If $n\gg M,N$, this can efficiently be solved by taking the dual.

- ***Non-linear classification***: In the simplest case, we can use a linearly parameterized family of functions $f(\boldsymbol{z}) = \boldsymbol{\theta} \cdot \boldsymbol{F}(\boldsymbol{z})$. Our problem then reduces to solving the following system of linear inequalities: $\boldsymbol{\theta} \cdot \boldsymbol{F}(\boldsymbol{x}^i) \geq 1 \ \forall i$ and $\boldsymbol{\theta} \cdot \boldsymbol{F}(\boldsymbol{y}^i) \leq -1 \ \forall i$.

Infinite-dimensional optimization

- *Hilbert spaces*
  - ***Definition***: A *pre-Hilbert Space* consists of a vector space V over $\mathbb{C}$ with an inner product $\langle \boldsymbol{x}, \boldsymbol{y} \rangle : V \times V \to \mathbb{C}$. Satisfies (a) $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \overline{\langle \boldsymbol{y}, \boldsymbol{x} \rangle}$ (b) $\langle \boldsymbol{x} + \boldsymbol{y}, \boldsymbol{z} \rangle = \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \langle \boldsymbol{y}, \boldsymbol{z} \rangle$ (c) $\langle \lambda \boldsymbol{x}, \boldsymbol{y} \rangle = \lambda \langle \boldsymbol{x}, \boldsymbol{y} \rangle$ for *all* λ (d) $\langle \boldsymbol{x}, \boldsymbol{x} \rangle \geq 0$, with equality iff $\boldsymbol{x} = \boldsymbol{0}$. $\|\boldsymbol{x}\| = \sqrt{\langle \boldsymbol{x}, \boldsymbol{x} \rangle}$ is a norm, and inner product is continuous in both its arguments under that norm (Proof on Luenberger, pp49).
    - V = space of sequences that are square-summable. Define $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \sum_{i=1}^{\infty} x_i y_i$. [Finite by Cauchy Schwarz].
    - $V = \mathcal{L}_2[a,b]$ = space of measurable functions $\boldsymbol{x} : [a,b] \to \mathbb{R}$ such that $|\boldsymbol{x}(t)|^2$ is Lesbegue integrable. $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \int_a^b \boldsymbol{x}(t) \boldsymbol{y}(t) \ \mathrm{d}t$.
    - V = polynomials on $[a, b]$ with $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \int_a^b \boldsymbol{x}(t) \boldsymbol{y}(t) \ \mathrm{d}t$
  - ***Definition***: A *Sequence* $\{x_n\} \subseteq V$ is *Cauchy* iff $\|\boldsymbol{x}_n - \boldsymbol{x}_m\| \to 0$ as $n, m \to \infty$, under the norm $\|\boldsymbol{y}\| = \sqrt{\langle \boldsymbol{y}, \boldsymbol{y} \rangle}$.
  - ***Definition***: A *Hilbert Space* is a *complete pre-Hilbert space*; one in which every Cauchy sequences converges in the space.
    - ***Example of an incomplete space***: Take $V = C[0,1]$, the space of continuous functions on [0,1]. Consider two norms

$$\|f\| = \max_{x \in [0,1]} |f(x)| \qquad \qquad \|f\|_* = \int_0^1 |f(x)| \ \mathrm{d}x$$

      Neither, it turns out, induce an inner product, so this it not a Hilbert Space. However:
      - V is complete under the first norm. Even though there is a sequence of functions f_n that tend to a step function (which isn't in

the space), the sequence isn't Cauchy, because $\left\| f_n - \text{step} \right\| \to \frac{1}{2}$, since we are considering the point of *maximum* difference.

- V is *not* complete under the starred norm, because $\left\| f_n - \text{step} \right\|_* \to 0$, so the sequence is Cauchy and leaves the space. Indeed, the area between f_n and the step function shrinks to 0. □

- ***Definition (Orthogonality)***: We say $\boldsymbol{x}$ and $\boldsymbol{y}$ are orthogonal iff $\left\langle \boldsymbol{x}, \boldsymbol{y} \right\rangle = 0$. We further define $M^\perp = \left\{ \boldsymbol{y} : \left\langle \boldsymbol{y}, \boldsymbol{x} \right\rangle = 0 \ \forall \boldsymbol{x} \in M \right\}$. By the joint continuity of the inner product, this is always closed.
- ***Theorem (Projection)***: Let H be a Hilbert Space and K a closed and non-empty convex subset of *H*. Then, for any $\boldsymbol{x} \in H$, $\min_{k \in K} \left\| \boldsymbol{x} - \boldsymbol{k} \right\|$ has an optimal solution $\boldsymbol{k}_0$ called the *projection* of $\boldsymbol{x}$ onto *K*. $\boldsymbol{k}' \in K$ is equal to $\boldsymbol{k}_0$ if and only if $\boldsymbol{x} - \boldsymbol{k}' \in K^\perp$ (in other words, $\left\langle \boldsymbol{x} - \boldsymbol{k}', \boldsymbol{k} - \boldsymbol{k}' \right\rangle \le 0 \ \forall \boldsymbol{k} \in K$).

  Proof: Luenberger, pp. 51.
- ***Theorem***: If M is a closed subspace of *H*, then $H = M \oplus M^\perp$ and $M = M^{\perp\perp}$. As such, we call $\boldsymbol{k}_0$ in the projection theorem the *orthogonal projection* of $\boldsymbol{x}$ onto *M*, and we can write $\boldsymbol{x} = \boldsymbol{k}_0 + \left(\boldsymbol{x} - \boldsymbol{k}_0 \right)$, where $\boldsymbol{k}_0 \in M, (\boldsymbol{x} - \boldsymbol{k}_0) \in M^\perp$.

  Proof: Luenberger, pp. 53.
- ***Definition (Linear functional)***: A function $\varphi : V \to \mathbb{C}$ is a *linear functional* if $\varphi(\alpha \boldsymbol{x} + \beta \boldsymbol{y}) = \alpha \varphi(\boldsymbol{x}) + \beta \varphi(\boldsymbol{y})$.
  - ***Continuous*** if $\left| \varphi(\boldsymbol{y}) - \varphi(\boldsymbol{x}) \right| \le \varepsilon \ \forall \boldsymbol{y} : \left\| \boldsymbol{x} - \boldsymbol{y} \right\| \le \delta$. If φ is continuous at $\boldsymbol{x}_0$, it is continuous everywhere. ***Proof***: $\left| \varphi(\boldsymbol{y}) - \varphi(\boldsymbol{x}) \right| = \left| \varphi(\boldsymbol{y} - \boldsymbol{x} + \boldsymbol{x}_0) - \varphi(\boldsymbol{x}_0) \right|$.
  - ***Bounded*** if $\exists M$ s.t. $\left| \varphi(\boldsymbol{y}) \right| \le M \left\| \boldsymbol{y} \right\| \ \forall \boldsymbol{y}$. Define norm $\boxed{\left\| \varphi \right\| = \inf \left\{ M : \left| \varphi(\boldsymbol{y}) \right| \le M \left\| \boldsymbol{y} \right\| \right\} = \sup_{\left\| \boldsymbol{y} \right\| = 1} \left| \varphi(\boldsymbol{y}) \right|}$ (last step by linearity).

  Notes:
  - If continuous, continuous at 0, $\left| \varphi(\boldsymbol{y}) \right| \le 1 \ \forall \left\| \boldsymbol{y} \right\| \le \delta$. As such, $\left| \varphi(\boldsymbol{z}) \right| = \left| \frac{\|\boldsymbol{z}\|}{\delta} \varphi\left(\delta \frac{\boldsymbol{z}}{\|\boldsymbol{z}\|} \right) \right| = \frac{\|\boldsymbol{z}\|}{\delta} \varphi\left(\delta \frac{\boldsymbol{z}}{\|\boldsymbol{z}\|} \right) \le \frac{\|\boldsymbol{z}\|}{\delta}$, since bit inside brackets has norm δ.

- If bounded, $|\varphi(\boldsymbol{z})| \le M\|\boldsymbol{z}\| \ \forall \boldsymbol{z}$ and so $|\varphi(\boldsymbol{z})| \le \varepsilon \ \forall \boldsymbol{z} : \|\boldsymbol{z}\| \le \frac{\varepsilon}{M}$. So continuous at 0, and therefore everywhere.
- ***Example of a non-bounded linear functional***: Let V be the space of all sequences with finitely many non-zero elements, with norm $\|\boldsymbol{x}\| = \max_k |x_k|$. Then $\varphi(\boldsymbol{x}) = \max_k |kx_k|$ is unbounded because we can push the non-zero elements of $\boldsymbol{x}$ to infinity without changing the norm but making the functional grow to infinity. □

- ***Theorem (Riesz-Frechet)***: If $\varphi(\boldsymbol{x})$ is a continuous linear functional, then there exists a $\boldsymbol{z} \in H$ such that $\varphi(\boldsymbol{x}) = \langle \boldsymbol{x}, \boldsymbol{z} \rangle$

  Proof: Let $M = \{\boldsymbol{y} : \varphi(\boldsymbol{y}) = 0\}$. Since the functional is continuous, M is closed. If $M = H$, set $\boldsymbol{z} = \boldsymbol{0}$. Else, choose $\boldsymbol{\gamma} \in M^\perp$.

  $$\begin{aligned} \varphi\left(\boldsymbol{x} - \tfrac{\varphi(\boldsymbol{x})}{\varphi(\boldsymbol{\gamma})}\boldsymbol{\gamma}\right) = \varphi(\boldsymbol{x}) - \varphi(\boldsymbol{x}) = 0 &\Rightarrow \boldsymbol{x} - \tfrac{\varphi(\boldsymbol{x})}{\varphi(\boldsymbol{\gamma})}\boldsymbol{\gamma} \in M \\ &\Rightarrow 0 = \left\langle \boldsymbol{x} - \tfrac{\varphi(\boldsymbol{x})}{\varphi(\boldsymbol{\gamma})}\boldsymbol{\gamma}, \boldsymbol{\gamma} \right\rangle = \langle \boldsymbol{x}, \boldsymbol{\gamma} \rangle - \tfrac{\varphi(\boldsymbol{x})}{\varphi(\boldsymbol{\gamma})}\langle \boldsymbol{\gamma}, \boldsymbol{\gamma} \rangle \\ &\Rightarrow \varphi(\boldsymbol{x}) = \left\langle \boldsymbol{x}, \tfrac{\overline{\varphi(\boldsymbol{\gamma})}}{\|\boldsymbol{\gamma}\|^2}\boldsymbol{\gamma} \right\rangle = \langle \boldsymbol{x}, \boldsymbol{z} \rangle \end{aligned}$$

  Note also that by Cauchy-Schwarz, $|\varphi(\boldsymbol{x})| \le \|\boldsymbol{z}\|\|\boldsymbol{x}\| \Rightarrow \|\varphi\| = \|\boldsymbol{z}\|$. ■

- This means that Hilbert spaces are self-dual (see later), and that we can write $\varphi(\boldsymbol{x}) = \langle \boldsymbol{x}, \varphi \rangle$.
- ***Theorem (Special case of the Hahn-Banach Theorem)***: Let $M \subseteq H$ be a closed subspace and φ_M be a continuous linear functional on M. Then there exists a continuous linear functional φ on H such that $\varphi(\boldsymbol{x}) = \varphi_M(\boldsymbol{x}) \ \forall \boldsymbol{x} \in M$ and $\|\varphi\| = \|\varphi_M\|$.

  Proof: Easy in the case of a Hilbert space. Since M is closed, it is also a Hilbert space, and so $\exists \boldsymbol{m} \in M$ such that $\varphi_M(\boldsymbol{x}) = \langle \boldsymbol{x}, \boldsymbol{m} \rangle$. Then define $\varphi(\boldsymbol{x}) = \langle \boldsymbol{x}, \boldsymbol{m} \rangle$ for $\boldsymbol{x} \in H$. By the CS inequality, $\|\varphi_M\| = \|\varphi\| = \|\boldsymbol{m}\|$. ■

- ***Banach Spaces & Their Duals***
  - A Banach space is a normed, complete vector space with *no* inner product.
    - $C[0,1]$ is the space of continuous function son [0,1], with

    $$\|\boldsymbol{f}\| = \max_{0 \le t \le 1} |\boldsymbol{f}(t)|$$

[As we showed above, the choice of this norm ensures completeness]. An example of a linear functional on this space is

$$\varphi(\boldsymbol{f}) = \int_0^1 \boldsymbol{x}(t)\ \mathrm{d}v(t) \le \left\|\boldsymbol{x}\right\| \int_0^1 \mathrm{d}v(t) \le \left\|\boldsymbol{x}\right\| \mathrm{TV}(v)$$

Provided the total variation of v, $\mathrm{TV}(v) < \infty$, where

$$\mathrm{TV}(v) = \sup_{\text{All partitions } 0=t_1<t_2<\cdots<t_n=1} \sum_{i=1}^n \left|v(t_i) - v(t_{i-1})\right|$$

- $\ell_p = \left\{\boldsymbol{x} \in \mathbb{R}^\infty : \left\|\boldsymbol{x}\right\|_p < \infty\right\}$, where

$$\left\|\boldsymbol{x}\right\|_p = \left(\sum_{i=1}^\infty \left|x_i\right|^p\right)^{1/p} \qquad \text{or} \quad \left\|\boldsymbol{x}\right\|_p = \sup_i \left|x_i\right| \ \text{if}\ p = \infty$$

- $\mathcal{L}_p[0,1] = \left\{\boldsymbol{x} : \int_0^1 \left|\boldsymbol{x}(t)\right|^p\ \mathrm{d}t < \infty\right\}$, with

$$\left\|\boldsymbol{f}\right\|_p = \left(\int_0^1 \left|\boldsymbol{f}(t)\right|^p\ \mathrm{d}t\right)^{1/p} \qquad 1 \le p < \infty$$ □

- ***Definition***: We say $V^* = \left\{\varphi : \varphi \text{ is continuous linear functional on } V\right\}$ is the *dual space* of V, with norm $\left\|\varphi\right\|_* = \sup\left\{\left|\varphi(\boldsymbol{x})\right| : \left\|\boldsymbol{x}\right\| \le 1\right\}$. $\left(V^*, \left\|\ \right\|_*\right)$ is always a Banach space.

  Proof: Want to show that $\forall \left\{\boldsymbol{x}_n^*\right\} \subseteq V^*$ with $\left\|\boldsymbol{x}_n^* - \boldsymbol{x}_m^*\right\|_* \le \varepsilon\ \forall n,m \ge M_\varepsilon$ converges to a point $\boldsymbol{x}^* = \lim_{n\to\infty} \boldsymbol{x}_n^* \in V^*$. First fix $\boldsymbol{x} \in V$ and note that

$$\left|\boldsymbol{x}_n^*(\boldsymbol{x}) - \boldsymbol{x}_m^*(\boldsymbol{x})\right| = \left|(\boldsymbol{x}_n^* - \boldsymbol{x}_m^*)(\boldsymbol{x})\right| \le \left\|\boldsymbol{x}_n^* - \boldsymbol{x}_m^*\right\|_* \left\|\boldsymbol{x}\right\|$$

  As such, $\left\{\boldsymbol{x}_n^*(\boldsymbol{x})\right\}$ is a Cauchy sequence in $\mathbb{R}$. Since $\mathbb{R}$ is complete, $\boldsymbol{x}^*(\boldsymbol{x}) = \lim_{n\to\infty} \boldsymbol{x}_n^*(\boldsymbol{x})$ exists. Define $\boldsymbol{x}^*$ pointwise using this limit. Now

  - ***Linearity***: By linearity of expectations, $\boldsymbol{x}^*$ is linear.
  - ***Continuity/boundedness***: Fix m_0 such that $\left\|\boldsymbol{x}_n^* - \boldsymbol{x}_m^*\right\| \le \varepsilon\ \forall n,m \ge m_0$. Then by the definition of $\boldsymbol{x}^*(\boldsymbol{x})$, $\left|\boldsymbol{x}^*(\boldsymbol{x}) - \boldsymbol{x}_m^*(\boldsymbol{x})\right| \le \varepsilon\left\|\boldsymbol{x}\right\|$, and

$$\left|\boldsymbol{x}^*(\boldsymbol{x})\right| \le \left|\boldsymbol{x}^*(\boldsymbol{x}) - \boldsymbol{x}_{m_0}^*(\boldsymbol{x})\right| + \left|\boldsymbol{x}_{m_0}^*(\boldsymbol{x})\right| \le \left(\varepsilon + \left\|\boldsymbol{x}_{m_0}^*\right\|_*\right)\left\|\boldsymbol{x}\right\| \Rightarrow \text{bounded}$$ ■

  Examples

  - We have already shown (Riesz-Frechet Theorem) that Hilbert spaces are self-dual.

- ***Theorem***: For $p \in (1,\infty)$, $\ell_p^* = \ell_q$, where $\frac{1}{p} + \frac{1}{q} = 1$. In other words

    1. For any $\boldsymbol{y} \in \ell_q$, $\boldsymbol{f}(\boldsymbol{x}) = \sum_{i=1}^{\infty} y_i x_i$ is a bounded linear functional on ℓ_p.
    2. Every $\boldsymbol{f} \in \ell_p^*$ can be represented uniquely as $\boldsymbol{f}(\boldsymbol{x}) = \sum_{i=1}^{\infty} y_i x_i$ with $\boldsymbol{y} \in \ell_q$.
    3. In both cases above, $\left\|\boldsymbol{f}\right\|_* = \left\|\boldsymbol{y}\right\|_q$.

    Proof: We prove each step separately

    1. Suppose $\boldsymbol{y} \in \ell_q$. Clearly, the proposed functional is linear. Furthermore, by Holder's Inequality $\left|\boldsymbol{f}(\boldsymbol{x})\right| \le \sum_{i=1}^{\infty} \left|y_i x_i\right| \le \left\|\boldsymbol{x}\right\|_p \left\|\boldsymbol{y}\right\|_q$. As such, the functional is also bounded, with $\left\|\boldsymbol{f}\right\| \le \left\|\boldsymbol{y}\right\|_q$.
    2. We prove this in four steps. Let $\boldsymbol{x} \in \ell_p$ and $\boldsymbol{f} \in \ell_p^*$.

        - ***Step 1*** *– approximate* $\boldsymbol{x}$ *using "basis" functions.* Define "basis functions" $\boldsymbol{e}^i \in \ell_p$, consisting of sequences which are identically 0 except for the i^{th} component. We then have

        $$\boldsymbol{x}_N = \sum_{i=1}^{N} x_i \boldsymbol{e}^i \to \boldsymbol{x}$$

        - ***Step 2*** *– find the image under* $\boldsymbol{f}$. In this case, let $y^i = \boldsymbol{f}(\boldsymbol{e}^i) \in \mathbb{R}$. We then have, by linearity of f

        $$f(\boldsymbol{x}_N) = \sum_{i=1}^{N} y^i x_i \to \sum_{i=1}^{\infty} y^i x_i$$

        By *continuity* of *f*, we have $f(\boldsymbol{x}_N) \to f(\boldsymbol{x})$. Thus

        $$f(\boldsymbol{x}) = \sum_{i=1}^{\infty} y^i x_i$$

        - ***Step 3*** *– show that the* y^i *form a vector in* ℓ_q. Define the vector $\boldsymbol{y}_N \in \ell_p$ as a "truncated y" series.

        $$\left(y_N\right)_i = \begin{cases} \left|y^i\right|^{q/p} \operatorname{sgn}\left(y^i\right) & i \le N \\ 0 & i > N \end{cases}$$

        We then have

$$\left\|\boldsymbol{y}_N\right\|_p = \left(\sum_{i=1}^N \left|y^i\right|^q\right)^{1/p}$$

$$f(\boldsymbol{y}_N) = \sum_{i=1}^N \left|y^i\right|^{q/p} \operatorname{sgn}(y^i) y^i = \sum_{i=1}^N \left|y^i\right|^q$$

We know that $\left|f(\boldsymbol{y}_N)\right| \leq \left\|\boldsymbol{f}\right\|_* \left\|\boldsymbol{y}_N\right\|_p \Rightarrow \frac{\left|f(\boldsymbol{y}_N)\right|}{\left\|\boldsymbol{y}_N\right\|_p} \leq \left\|\boldsymbol{f}\right\|_*$. Thus

$$\left(\sum_{i=1}^N \left|y^i\right|^q\right)^{1-(1/p)} = \left(\sum_{i=1}^N \left|y^i\right|^q\right)^{1/q} \leq \left\|\boldsymbol{f}\right\|_* \qquad \forall N$$

As such, $\left\|\boldsymbol{y}\right\|_q \leq \left\|\boldsymbol{f}\right\|$, and $\boldsymbol{y} \in \ell_q$.

3. We have shown that $\left\|\boldsymbol{y}\right\|_q \leq \left\|\boldsymbol{f}\right\|$ and $\left\|\boldsymbol{y}\right\|_q \geq \left\|\boldsymbol{f}\right\|$. Thus, $\left\|\boldsymbol{y}\right\|_q = \left\|\boldsymbol{f}\right\|$.

- ***Theorem***: $\ell_1^* = \ell_\infty$.

  Proof: Show this as above, and define the truncated vector $\boldsymbol{y}_N$ as a containing $\operatorname{sgn} y^N$ for $i = N$.

- ***Theorem***: For $p \in [1,\infty)$, $L_p^*[0,1] = L_q[0,1]$, where $p^{-1} + q^{-1} = 1$. For every $\boldsymbol{f} \in L_p^*$, there exists a $\boldsymbol{y} \in L_q$ such that $\boldsymbol{f}(\boldsymbol{x}) = \int_0^1 \boldsymbol{x}(t)\boldsymbol{y}(t)\ \mathrm{d}t$ and $\left\|\boldsymbol{f}\right\|_* = \left\|\boldsymbol{y}\right\|_q$.

- ***Theorem***: $C[0,1]^* = \mathrm{NBV}[0,1]$; we will prove this later using the HB Theorem.

- ***The Hahn Banach Theorem & Application to C[0,1]***
  - ***Theorem (Hahn-Banach)***: Let p be a continuous seminorm (same as a norm, except for the fact it can be equal to 0 even when $\boldsymbol{x} \neq \boldsymbol{0}$), $M \subseteq V$ be a closed subspace and $f : M \to \mathbb{R}$ be a linear functional such that $f(\boldsymbol{x}) \leq p(\boldsymbol{x})\ \forall \boldsymbol{x} \in M$. **Then**, there exists a linear functional $F : V \to \mathbb{R}$ such that $F(\boldsymbol{x}) \leq p(\boldsymbol{x})\ \forall \boldsymbol{x} \in V$ and $F(\boldsymbol{x}) = f(\boldsymbol{x})\ \forall \boldsymbol{x} \in M$.

    Note: Consider setting $p(\boldsymbol{x}) = \left\|\boldsymbol{f}\right\|_* \left\|\boldsymbol{x}\right\|$. Clearly, the condition of the theorem then applies, because $f(\boldsymbol{x}) \leq \left\|\boldsymbol{f}\right\|_* \left\|\boldsymbol{x}\right\|$. The theorem then implies that $\left\|\boldsymbol{F}\right\|_* = \left\|\boldsymbol{f}\right\|_*$.

The only reason we generalize p to a seminorm is to prove the geometric HB theorem (see later).

Note 2: The idea it is possible to extend $\boldsymbol{f}$ over an entire space is not particularly revolutionary. The crux of the theorem is that this extension has bounded norm. In a way, the HB Theorem can be stated as "the optimization problem $\min_{\boldsymbol{F}\in X^*} \left\|\boldsymbol{F}\right\|_*$ s.t. $\left\langle \boldsymbol{x}, \boldsymbol{F}\right\rangle = \boldsymbol{f}(\boldsymbol{x}) \ \ \forall \boldsymbol{x} \in M \subset X$ has a global optimium, and its value is $\left\|\boldsymbol{f}\right\|_*$".

Note 3: Let X be a normed vector space. Then $\forall \boldsymbol{x} \in X, \exists \boldsymbol{F}$ s.t. $\boldsymbol{F}(\boldsymbol{x}) = \left\|\boldsymbol{F}\right\|\left\|\boldsymbol{x}\right\|$. Define $\boldsymbol{f}(\alpha \boldsymbol{x}) = a\left\|\boldsymbol{x}\right\|$ on the subspace generated by $\boldsymbol{x}$; this has norm unity. By the HB Theorem, we can extend this to $\boldsymbol{F}$ on X with norm unity. This satisfies the requirements for the point $\boldsymbol{x}$.

- ***Example/Theorem (dual space of C*[0, 1]*, Riesz Representation Theorem)***: (In all that follows, use the usual norm, $\max_{0\le t\le 1}\left|\boldsymbol{x}(t)\right|$, over $C[0,1]$)
  - Take any function v of bounded variation on $[0,1]$. Then $\boldsymbol{f} : C[0,1] \to \mathbb{R}$ defined by $\boldsymbol{f}(\boldsymbol{x}) = \int_0^1 \boldsymbol{x}(t)\ \mathrm{d}v(t)$ is a bounded linear functional in $C[0,1]^*$.
  - Tale any bounded linear functional $\boldsymbol{f} \in C[0,1]^*$. Then there is a function v of bounded variation on [0, 1] such that $\boldsymbol{f}(\boldsymbol{x}) = \int_0^1 \boldsymbol{x}(t)\ \mathrm{d}v(t)$.
  - For the function defined in (2), $\left\|\boldsymbol{f}\right\| = \mathrm{TV}(v)$.
  - $C[0,1]^* = \mathrm{NBV}$.

Proof:

1. Clearly, any $\boldsymbol{f}$ defined in this fashion is linear. Furthermore, it is bounded:
$$\boldsymbol{f}(\boldsymbol{x}) = \int_0^1 \boldsymbol{x}(t)\ \mathrm{d}v(t) \le \max_{0\le t\le 1}\left|\boldsymbol{x}(t)\right|\mathrm{TV}(v) \le \left\|\boldsymbol{x}\right\|\mathrm{TV}(v)$$
2. Note that $C[0,1]$ (space of continuous functions on $[0,1]$) is a subset of $B[0,1]$, the space of bounded functions on $[0,1]$. Thus, by the HB Theorem
$$\boldsymbol{f} \in C[0,1]^* \Rightarrow \exists \boldsymbol{F} \in B[0,1]^* \text{ s.t. } \left\|\boldsymbol{F}\right\| = \left\|\boldsymbol{f}\right\|$$
   Our proof then goes as follows:

- ***Step 1*** *– Approximate $\boldsymbol{x} \in C[0,1]$ by discretisng it.* Define a set of step functions $\boldsymbol{u}_s(t) = \mathbb{1}_{\{t \le s\}} \in B[0,1]$. (Note these are not in $C[0,1]$; this is why it's useful to move to the larger space). We can write

$$\boldsymbol{x}(\tau) \approx \boldsymbol{z}^\pi(\tau) = \sum_{i=1}^n \boldsymbol{x}(t_i)\left(\boldsymbol{u}_{t_i}(\tau) - \boldsymbol{u}_{t_{i-1}}(\tau)\right) \in B[0,1]$$

Where π is some partition of $[0,1]$. Hereafter, when we write "tends to", we mean "as the partition gets arbitrarily fine".

- ***Step 2*** *– find the image under $\boldsymbol{F}$.* Let $v(s) = \boldsymbol{F}(\boldsymbol{u}_s) \in \mathbb{R}$ be the image of these "basis functions". Since $\boldsymbol{F}$ is linear and the first term in the sum is a constant, we can write

$$\boldsymbol{F}(\boldsymbol{z}^\pi) = \sum_{i=1}^n \boldsymbol{x}(t_i)\left(v(t_i) - v(t_{i-1})\right) \to \int_0^1 \boldsymbol{x}(t) \ \mathrm{d}v(t) \in \mathbb{R}$$

- ***Step 3*** *– bridge the gap between B and C.* By uniform continuity of $\boldsymbol{x}$, the approximation $\boldsymbol{z}^\pi$ becomes arbitrarily good (using the max norm). Since $\boldsymbol{F}$ is also continuous using the max norm, this implies that $\boldsymbol{F}(\boldsymbol{z}^\pi) \to \boldsymbol{F}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x})$. As such

$$\boldsymbol{f}(\boldsymbol{x}) = \int_0^1 \boldsymbol{x}(t) \ \mathrm{d}v(t)$$

- ***Step 4*** *– show v has bounded TV.* By linearity of $\boldsymbol{F}$, we have

$$\begin{aligned}\sum_{i=1}^n \left|v(t_i) - v(t_{i-1})\right| &= \boldsymbol{F}\left(\sum_{i=1}^n \varepsilon_i \left(u_{t_i} - u_{t_{i-1}}\right)\right) \\ &\le \left\|\boldsymbol{F}\right\|_* \left\|\sum_{i=1}^n \varepsilon_i \left(u_{t_i} - u_{t_{i-1}}\right)\right\| \\ &= \left\|\boldsymbol{F}\right\|_* = \left\|\boldsymbol{f}\right\|_*\end{aligned}$$

Where $\varepsilon_i = \pm 1$ takes the absolute value into account, and the jump from line 2 to 3 follows since u are step functions. Taking a supremum over all partitions, we find $\mathrm{TV}(v) \le \left\|\boldsymbol{f}\right\|_* < \infty$.

Note, however, that the F produced by the HB Theorem is not necessarily unique. As such, nor is the function v. This theorem only states there exists such a v.

3. From (1), it is clear that $\left\|\boldsymbol{f}\right\|_* \le \mathrm{TV}(v)$. From (2, Step 4) it is clear that $\left\|\boldsymbol{f}\right\|_* \ge \mathrm{TV}(v)$. Thus, $\left\|\boldsymbol{f}\right\|_* = \mathrm{TV}(v)$.

4. Because if not non-uniqueness noted above, $C[0,1]^* \neq \mathrm{BV}[0,1]$. Indeed, the linear functional $f(\boldsymbol{x}) = \boldsymbol{x}(\frac{1}{2})$, for example, we can represented using a v that is 0 on [0, ½), 1 on (½, 1] and takes *any* value at ½. As such, we define the space NBV[0,1] (*normalized bounded variation*), consisting of all functions of bounded variations that vanish at 0 and are right-continuous. We then have $C[0,1]^* = \mathrm{NBV}[0,1]$, because every element in one set can be mapped to an element in the other. ■

- ***Odds & Ends***
  - We will sometimes abuse notation and write $\boldsymbol{x}^*(\boldsymbol{x}) = \left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle$, for $\boldsymbol{x} \in V, \boldsymbol{x}^* \in V^*$. In a Hilbert space, this is true, because $V = V^*$ (Riesz-Frechet). In a Banach space, it is convenient notation (see Hyperplanes section).
  - This means that by viewing $\boldsymbol{x}$ as fixed, $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle$ also defines a functional in X^{**} (easy to show linear and bounded). Now, we have
    - $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle \leq \left\| \boldsymbol{x} \right\| \left\| \boldsymbol{x}^* \right\|$
    - By Corollary 2 of H-B (below), $\forall \boldsymbol{x} \in X, \exists \boldsymbol{x}^* \in X^*$ s.t. $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle = \left\| \boldsymbol{x} \right\| \left\| \boldsymbol{x}^* \right\|$

    As such, if we consider $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle$ as a functional on X^{**}, its norm is $\left\| \boldsymbol{x} \right\|_X$. We define the *natural mapping* $\varphi : X \to X^{**}$ so that $\varphi(\boldsymbol{x})$ maps $\boldsymbol{x}$ to the functional it generates in X^{**}; in other words, $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle = \left\langle \varphi(\boldsymbol{x}), \boldsymbol{x}^* \right\rangle$ but with $\varphi(\boldsymbol{x}) \in \boldsymbol{x}^{**}$. The mapping is linear, and, as we showed previously, norm-preserving; $\left\| \varphi(\boldsymbol{x}) \right\|_{X^{**}} = \left\| \boldsymbol{x} \right\|_X$. But it might not be onto – some elements in X^{**} might not be representable by elements in X. If $X = X^{**}$, X is called **reflexive**. All Hilbert spaces are reflexive, as are ℓ_p and L_p for $p \in (1, \infty)$.
  - We have $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle \leq \left\| \boldsymbol{x}^* \right\| \left\| \boldsymbol{x} \right\|$. In a Hilbert space, we have equality if and only if $\boldsymbol{x}^* = \alpha \boldsymbol{x}$. In a Banach space, we say $\boldsymbol{x}^* \in X^*$ is *aligned* with $\boldsymbol{x} \in X$ if $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle = \left\| \boldsymbol{x}^* \right\| \left\| \boldsymbol{x} \right\|$. They are said to be *orthogonal* if $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle = 0$. Similarly, if $S \subseteq X$, we say $S^\perp = \left\{ \boldsymbol{x}^* \in X^* : \left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle = 0 \ \forall \boldsymbol{x} \in S \right\} \subseteq X^*$. If $U \subseteq X^*$, we say $U^\perp \cap X = {}^\perp S^* = \left\{ \boldsymbol{x} \in X : \left\langle \boldsymbol{x}^*, \boldsymbol{x} \right\rangle = 0 \ \forall \boldsymbol{x}^* \in S^* \right\} \subseteq X$.

- **Theorem**: If $M \subseteq X$, then ${}^{\perp}(M^{\perp}) = M$.

  Proof: Clearly, $M \subseteq {}^{\perp}(M^{\perp})$. To show the converse, we'll show that $\boldsymbol{x} \notin M \Rightarrow \boldsymbol{x} \notin {}^{\perp}(M^{\perp})$. Define a linear functional $\boldsymbol{f}$ on the space spanned by M and $\boldsymbol{x}$ which vanishes on M so that $\boldsymbol{f}(\boldsymbol{m} + \alpha\boldsymbol{x}) = \alpha$. It can be shown that $\left\|\boldsymbol{f}\right\| < \infty$, and so by the HB Theorem, we can extend it to some $\boldsymbol{F}$ which also vanishes on M. As such, $\boldsymbol{F} \in M^{\perp}$. However, $\boldsymbol{F}(\boldsymbol{x}) = \left\langle \boldsymbol{F}, \boldsymbol{x} \right\rangle = 1 \neq 0$, and so $\boldsymbol{x} \notin {}^{\perp}(M^{\perp})$. ■

Minimum Norm Problems

- Let us consider a vector $\boldsymbol{x} \in X$. There are clearly two ways to take the norm of that vector – as an element of X or as an element of X^{**} (a functional on X^{*}).

$$\left\|\boldsymbol{x}\right\| \qquad \text{or} \qquad \max_{\left\|\boldsymbol{x}^*\right\|=1} \left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle$$

  It is clear these two should be equal, because $\left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle \leq \left\|\boldsymbol{x}\right\| \cdot 1$ (or, more intuitively, because the second norm finds the most $\boldsymbol{x}$ can yield under a functional of norm 1 – clearly, the answer is its norm). Let us now restrict ourselves to a subspace M of X. We can, again, define two norms

$$\left\|\boldsymbol{x}\right\|_M = \inf_{\boldsymbol{m} \in M} \left\|\boldsymbol{x} - \boldsymbol{m}\right\| \qquad \text{or} \qquad \left\|\boldsymbol{x}\right\|_{M^{\perp}} = \sup_{\substack{\left\|\boldsymbol{x}^*\right\|=1 \\ \boldsymbol{x}^* \in M^{\perp}}} \left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle$$

  The first simply consists of the minimum distance between $\boldsymbol{x}$ and M (as opposed to between $\boldsymbol{x}$ and 0). The second is the most $\boldsymbol{x}$ can yield under a functional of norm 1 that annihilates any element of M. Intuitively, the "remaining bit" that's "not annihilated" is $\boldsymbol{x} - \boldsymbol{m}$; this is maximized when it is aligned with $\boldsymbol{x}^*$ – at $\boldsymbol{m}_0$. So it makes sense that the two should be equal.

- **Theorem**: Consider a normed linear space X and a subspace M therein. Let $\boldsymbol{x} \in X$. Then

$$d = \inf_{\boldsymbol{m} \in M} \left\|\boldsymbol{x} - \boldsymbol{m}\right\| = \max_{\substack{\left\|\boldsymbol{x}^*\right\|_* \leq 1 \\ \boldsymbol{x}^* \in M^{\perp}}} \left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle \qquad \left(= \max_{\substack{\left\|\boldsymbol{x}^*\right\|_* = 1 \\ \boldsymbol{x}^* \in M^{\perp}}} \left\langle \boldsymbol{x} - \boldsymbol{m}_0, \boldsymbol{x}^* \right\rangle \right).$$

  Or, in our terminology above, $\left\|\boldsymbol{x}\right\|_M = \left\|\boldsymbol{x}\right\|_{M^{\perp}}$. The maximum on the right is achieved for some $\boldsymbol{x}_0^* \in M^{\perp}$; if the infemum on the left is achieved for some $\boldsymbol{m}_0 \in M$, then $\boldsymbol{x} - \boldsymbol{m}_0$ is aligned with $\boldsymbol{x}_0^*$.

Intuitively, this is because at the optimal $\boldsymbol{m}$, the residual $\boldsymbol{x} - \boldsymbol{m}_0$ is aligned to some vector in $M^\perp$. As such, for that vector, $\langle \boldsymbol{x} - \boldsymbol{m}_0, \boldsymbol{x}^* \rangle = \|\boldsymbol{x} - \boldsymbol{m}_0\|$. For every other $\boldsymbol{x}^*$, it'll be smaller than that.

Pictorially, looking for the point on M that minimizes the norm is equivalent to looking for a point on $M^\perp$ that is aligned with $\boldsymbol{x} - \boldsymbol{m}_0$.

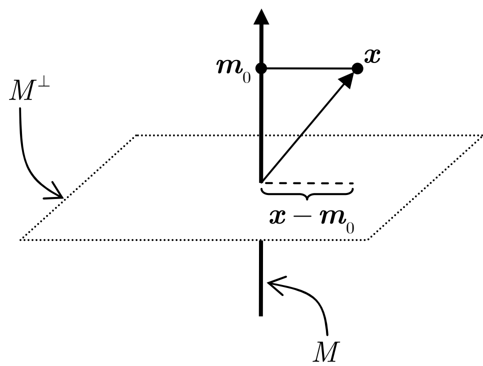

This also implies that a vector $\boldsymbol{m}_0$ is the minimum-norm projection if and only if there is a non-zero vector $\boldsymbol{x}^* \in M^\perp$ aligned with $\boldsymbol{x} - \boldsymbol{m}_0$.

- ***Theorem***: Let M be a subspace in a real normed space X. Let $\boldsymbol{x}^* \in X^*$. Then

$$d = \min_{\boldsymbol{m}^* \in M^\perp} \|\boldsymbol{x}^* - \boldsymbol{m}^*\| = \sup_{\boldsymbol{x} \in M, \|\boldsymbol{x}\| \le 1} \langle \boldsymbol{x}, \boldsymbol{x}^* \rangle$$

  where the minimum on the left is achieved for some $\boldsymbol{m}_0^* \in M^\perp$. If the supremum is achieved for some $\boldsymbol{x}_0 \in M$, then $\boldsymbol{x}^* - \boldsymbol{m}_0^*$ is aligned with $\boldsymbol{x}_0$.

  Because the minimum on the left is always achieved, it is *always* more desirable to express optimization problems in a dual space.

- In many optimization problems, we seek to minimize a norm over an affine subset of a dual space rather than subspace. More specifically, subject to a set of linear constraints of the form $\langle \boldsymbol{y}_i, \boldsymbol{x}^* \rangle = c_i$. In that case, if $\bar{\boldsymbol{x}}^*$ is some vector that satisfies these constraints,

$$\begin{aligned}\min_{\langle \boldsymbol{y}_i, \boldsymbol{x}^* \rangle = \boldsymbol{c}_i} \left\| \boldsymbol{x}^* \right\| &= \min_{\boldsymbol{m}^* \in M^\perp} \left\| \overline{\boldsymbol{x}}^* - \boldsymbol{m}^* \right\| \\ &= \sup_{\substack{\|\boldsymbol{x}\| \le 1 \\ \boldsymbol{x} \in M}} \left\langle \boldsymbol{x}, \overline{\boldsymbol{x}}^* \right\rangle \\ &= \sup_{\|\sum a_i \boldsymbol{y}_i\| \le 1} \left\langle \sum a_i \boldsymbol{y}_i, \overline{\boldsymbol{x}}^* \right\rangle \\ &= \sup_{\|\sum a_i \boldsymbol{y}_i\| \le 1} \boldsymbol{c} \cdot \boldsymbol{a}\end{aligned}$$

Where M is the space generated by the $\boldsymbol{y}_i$. The last equality follows from the fact that $\overline{\boldsymbol{x}}^*$ satisfies the equalities. Note that the optimal $\sum a_i y_i$ is aligned with the optimal $\boldsymbol{x}^*$.

- *Applications*
  - ***Example***: Consider the problem of selecting the field current $u(t)$ on $[0,1]$ to drive a motor from initial conditions $\theta(0) = \dot\theta(0) = 0$ to state $\theta(1) = 1$, $\dot\theta(1) = 0$ while minimizing $\max_{0 \le t \le 1} |u(t)|$. Assume the motor is governed by $\ddot\theta(t) + \dot\theta(t) = u(t)$.
    - First, we need to choose a space on which to optimize our problem. Choose, $u(t) \in L_\infty[0,1]$, which is the dual of $L_1[0,1]$.
    - First, note that we can treat the governing equation as a first-order equations and use an integrating factor to find

$$\begin{aligned}\tfrac{\mathrm{d}}{\mathrm{d}t}\left(e^t \dot\theta(t)\right) = e^t u(t) &\Rightarrow \left[e^t \dot\theta(t)\right]_0^1 = \int_0^1 e^t u(t) \\ &\Rightarrow \dot\theta(1) = \int_0^1 e^{t-1} u(t)\ \mathrm{d}t \Rightarrow \boxed{\dot\theta(1) = \left\langle e^{t-1}, u \right\rangle}\end{aligned}$$

(Here e^{t-1} is considered as a function in L_1 and $\int_0^1 e^{t-1} u(t)\ \mathrm{d}t$ as some functional in L_1^* on that function).

We can also integrate the governing equation directly to get

$$\dot\theta(1) - \dot\theta(0) + \theta(1) - \theta(0) = \int_0^1 u(t)\ \mathrm{d}t \Rightarrow \theta(1) = \int_0^1 u(t)\ \mathrm{d}t - \dot\theta(1)$$

Feeding in the results of the previous equation

$$\theta(1) = \int_0^1 \left(1 - e^{t-1}\right) u(t)\ \mathrm{d}t \Rightarrow \boxed{\theta(1) = \left\langle 1 - e^{t-1}, u \right\rangle}$$

    - As such, our problem boils down to minimizing the norm of u subject to $\left\langle e^{t-1}, u \right\rangle = 0$ and $\left\langle 1 - e^{t-1}, u \right\rangle = 1$. This is precisely the situation considered at the end of the previous section, and so this optimization problem is equivalent to

$$\max_{\left\|a_1 e^{t-1}+a_2\left(1-e^{t-1}\right)\right\|\leq 1}\left(1\cdot a_2+0\cdot a_1\right)$$

Where the norm is taken in L_1 (the primal space). As such, we want to maximize a_2 subject to $\int_0^1\left|\left(a_1-a_2\right)e^{t-1}+a_2\right|\mathrm{d}t\leq 1$.

- Once we have found the optimal value of a_2, we can find u by characterizing the alignment between L_1 and L_∞. For $x\in L_1$ and $u\in L_\infty$ to be aligned, we require

$$\left\langle x,u\right\rangle=\int_0^1 x(t)u(t)\ \mathrm{d}t=\left\|x\right\|\left\|u\right\|_*=\int_0^1\left|x(t)\right|\ \mathrm{d}t\cdot\max_{t\in[0,1]}\left|u(t)\right|$$

  For this to be true, it is clear that u can only take two values ($\pm M$) and that it must have the same sign as x at any given value of t.

- Finally, consider that in this case, x is $\left(a_1-a_2\right)e^{t-1}+a_2$. Clearly, it changes sign at most once. And so $u(t)$ must be equal to $\pm M$ with a single change in sign.

o ***Example***: Consider selecting a thrust program $u(t)$ for a vertically ascending rocket subject only to gravity and thrust in order to reach a given altitude (say 1) with minimum fuel expenditure $\int_0^T\left|u(t)\right|\mathrm{d}t$. Assume $x(0)=\dot{x}(0)=0$, unit mass and unit gravity. The equation of motion is $\ddot{x}(t)=u(t)-1$.

- We might originally regard this problem in L_1, but this is not a dual space. Instead, consider it in $NBV[0,1]$, and associate with every u a $v\in NBV[0,1]$ so that $\left|u(t)\right|\ \mathrm{d}t=\left|dv(t)\right|$.

- The *time* at which the rocket needs to reach an altitude of 1. We denote this by an unknown T and then optimize over this parameter. Integrating the equation of motion

$$\dot{x}(t)-\dot{x}(0)=\int_0^t u(s)\ \mathrm{d}s-t\Rightarrow\dot{x}(t)=\int_0^t u(s)\ \mathrm{d}s-t$$

  Integrating again, by parts

$$x(T) = \int_0^T \int_0^t u(s)\ \mathrm{d}s\ \mathrm{d}t - \frac{T^2}{2}$$
$$x(T) = \left[t\int_0^t u(s)\ \mathrm{d}s\right]_0^T - \int_0^T tu(t)\ \mathrm{d}t - \frac{T^2}{2}$$
$$x(T) = T\int_0^T u(s)\ \mathrm{d}s - \int_0^T tu(t)\ \mathrm{d}t - \frac{T^2}{2}$$
$$\int_0^T (T-t)u(s)\ \mathrm{d}s = x(T) + \frac{T^2}{2}$$
$$\boxed{\left\langle T-t, v\right\rangle = x(T) + \frac{T^2}{2}}$$

Where v is the function in $NBV[0,1]$ associated with u, as described above.

- Our problem is then a minimum norm problem subject to a single linear constraint. We want $x(T) = 1$, and using our theorem, as we did above

$$\min_{\left\langle T-t,v\right\rangle = 1+\frac{1}{2}T^2} \left\|v\right\| = \max_{\left\|(T-t)a\right\|\le 1}\left[a\left(1+\tfrac{1}{2}T^2\right)\right]$$

This is a one-dimensional problem. The norm is in $C[0,1]$, the space to which NBV is dual. As such $\left\|(T-t)a\right\| = \max_{t\in[0,1]}\left|(T-t)a\right| = T\left|a\right|$, and the optimum occurs at $a = 1/T$. We then have $\min_{\left\langle T-t,v\right\rangle = 1+\frac{1}{2}T^2}\left\|v\right\| = \frac{1}{T}+\frac{1}{2}T$. Differentiating this with respect to T, we find that the minimum fuel expenditure of $\sqrt{2}$ is achieved at $T = \sqrt{2}$.

- To find the optimal u, note that the optimal v must be aligned to $(T-t)a$. As we discussed above when characterizing alignment of C and NBV, this means that v must be a step function at $t = 0$, rising to $\sqrt{2}$ at $t = 0$, and as such, u must be an impulse (delta function) at $t = 0$.

- ***Hyperplanes & the Geometric Hahn-Banach Theorem***
  - ***Definition***: A *hyperplane* H of a normed linear space X is a maximal proper affine set. ie: if $H \subseteq A$ and A is affine, then either $A = H$ or $A = X$.
  - ***Theorem***: A set H is a hyperplane if and only if it is of the form $\left\{\boldsymbol{x}\in X : f(\boldsymbol{x}) = c\right\}$ where f is a non-zero linear functional, and c is a scalar.

    Proof:

    - ***If***: Let $H = \boldsymbol{x}_0 + M$, where M is a linear subspace

- If $\boldsymbol{x}_0 \notin M$, then $X = \boldsymbol{x}_0 + \text{span}\left(\boldsymbol{x}_0, M\right)$ by the maximality property of H, since this set is bigger than M. Thus, $X = \text{span}\left(\boldsymbol{x}_0, M\right)$. We can therefore write any $\boldsymbol{x} \in X$ as $\boldsymbol{x} = \alpha \boldsymbol{x}_0 + \boldsymbol{m}$, with $\boldsymbol{m} \in M$. Define $f(\alpha \boldsymbol{x}_0 + \boldsymbol{m}) = \alpha$. Then $H = \left\{\boldsymbol{x} : f(\boldsymbol{x}) = 1\right\}$.
- If $\boldsymbol{x}_0 \in M$, then $H = M$. Simply pick $\boldsymbol{x}_0' \notin M$, apply the above and get $H = \left\{\boldsymbol{x} : f(\boldsymbol{x}) = 0\right\}$.

- ***Only if***: Suppose $\boldsymbol{f} \neq \boldsymbol{0}$ and let $M = \left\{\boldsymbol{x} : \boldsymbol{f}(\boldsymbol{x}) = 0\right\}$. Clearly, it is a linear subspace. Furthermore, there exists $\boldsymbol{x}_0$ such that $f(\boldsymbol{x}_0) = 1$. As such, $\boldsymbol{x} - f(\boldsymbol{x})\boldsymbol{x}_0 \in M \ \forall \boldsymbol{x}$, and so $X = \left\{\boldsymbol{v} + f(\boldsymbol{x})\boldsymbol{x}_0 : \boldsymbol{v} \in M\right\}$. So we only require *one* extra vector ($\boldsymbol{x}_0$) to expand M into the whole subspace. So M is a maximal subspace. Thus $H = \left\{\boldsymbol{x} : \boldsymbol{f}(\boldsymbol{x}) = c\right\} = \left\{\boldsymbol{x} = c\boldsymbol{x}_0 + M\right\}$ is a hyperplane. ■

Important note: A hyperplane is only closed if *f* is linear and continuous.

- ***Theorem (Geometric Hahn-Banach)***: Let K be a convex set having a nonempty interior in a real normed linear vector space X. Suppose V is an affine set in X containing no interior points of K. Then there is a closed hyperplane in X containing V but containing to interior point of K. In other words, there is an element $\boldsymbol{x}^* \in X^*$ and a constant c such that $\left\langle \boldsymbol{v}, \boldsymbol{x}^* \right\rangle = c$ for all $\boldsymbol{v} \in V$ and $\left\langle \boldsymbol{k}, \boldsymbol{x}^* \right\rangle < c$ for all $\boldsymbol{k} \in K$.
- We will abuse notation and write $\boldsymbol{x}^*(\boldsymbol{x}) = \left\langle \boldsymbol{x}, \boldsymbol{x}^* \right\rangle$, for $\boldsymbol{x} \in V, \boldsymbol{x}^* \in V^*$.
  - In Hilbert Spaces, this is actually true thanks to Riesz-Frechet.
  - It allows to represent *all* hyperplanes as $\left\{\boldsymbol{x} : \left\langle \boldsymbol{x}, \boldsymbol{a} \right\rangle = 0\right\}$.

Mathematical background

- ***Vector spaces & Topology***
  - ***Inner products***

- *Cauchy-Schwarz inequality*: $\left|\langle \boldsymbol{x},\boldsymbol{y}\rangle\right|^2 \le \langle \boldsymbol{x},\boldsymbol{x}\rangle \cdot \langle \boldsymbol{y},\boldsymbol{y}\rangle$, with equality if and only if $\boldsymbol{x} = \lambda \boldsymbol{y}$, or either vectors are $\boldsymbol{0}$.

  Proof: If $\boldsymbol{y} = 0$, the result is simple. Else

  $$0 \le \langle \boldsymbol{x} - \lambda\boldsymbol{y}, \boldsymbol{x} - \lambda\boldsymbol{y}\rangle = \langle \boldsymbol{x},\boldsymbol{x}\rangle - 2\lambda\langle \boldsymbol{x},\boldsymbol{y}\rangle + \lambda^2\langle \boldsymbol{y},\boldsymbol{y}\rangle$$

  Set $\lambda = \langle \boldsymbol{x},\boldsymbol{y}\rangle / \langle \boldsymbol{y},\boldsymbol{y}\rangle$ to get the result. ■

- *Parallelogram Law*: $\|\boldsymbol{x}+\boldsymbol{y}\|^2 + \|\boldsymbol{x}-\boldsymbol{y}\|^2 = 2\|\boldsymbol{x}\|^2 + 2\|\boldsymbol{y}\|^2$ (prove by extending norms as inner products).

- Induced norm $\sqrt{\langle \boldsymbol{x},\boldsymbol{x}\rangle}$ satisfies triangle (expand $\|\boldsymbol{x}+\boldsymbol{y}\|^2$, and use C-S).

- For matrices, the standard inner product is $\langle X, Y\rangle = \operatorname{tr}(X^\top Y)$. Equivalent to multiplying every element in X with the corresponding element in Y. The induced norm is the Frobenius Norm, $\|X\|_F$.

- ***Norms***

  - A norm has the properties (a) $\|\boldsymbol{x}\| \ge 0$ (b) $\|\boldsymbol{x}\| = 0 \Leftrightarrow \boldsymbol{x} = \boldsymbol{0}$ (c) $\|\alpha\boldsymbol{x}\| = |\alpha|\|\boldsymbol{x}\|$ (d) $\|\boldsymbol{x}+\boldsymbol{y}\| \le \|\boldsymbol{x}\| + \|\boldsymbol{y}\|$. A seminorm might not satisfy (b).

  - ***Common norms***: $\|\boldsymbol{x}\|_p = \left(\sum | x_i |^p\right)^{1/p}$. For $P \in \mathbb{S}^n_{++}$, $\|\boldsymbol{x}\|_P = \sqrt{\boldsymbol{x}^\top P \boldsymbol{x}}$ (the unit ball is an ellipsoid). $\|\boldsymbol{x}\|_\infty = \max\left\{| x_1 |, \cdots | x_n |\right\}$.

  - The *dual norm* of a norm $\|\ \|$ is $\|\boldsymbol{z}\|_* = \sup_{\boldsymbol{x}}\left\{\boldsymbol{z}\cdot\boldsymbol{x} : \|\boldsymbol{x}\| \le 1\right\}$. It is the support function of the unit ball of the norm. Note that $\boldsymbol{x}\cdot\boldsymbol{y} \le \|\boldsymbol{x}\|\|\boldsymbol{y}\|_*$

    - For $p \in (1,\infty)$, the dual norm of $\|\ \|_p$ is $\|\ \|_q$, where $\frac{1}{p} + \frac{1}{q} = 1$.
    - $\|\ \|_{**} = \|\ \|$ (not true in infinite dimensional spaces).

- ***Open/closed sets***

  - $N_r(\boldsymbol{x}) = \left\{\boldsymbol{y} \in X : \|\boldsymbol{x}-\boldsymbol{y}\| < r\right\}$ is a *neighborhood* of $\boldsymbol{x}$ ("open ball")
  - $\boldsymbol{x} \in \operatorname{int}\mathcal{E}$ if $\exists r$ s.t. $N_r(\boldsymbol{x}) \subset \mathcal{E}$. $\mathcal{E}$ open $\Leftrightarrow \mathcal{E} = \operatorname{int}\mathcal{E}$.
  - $\boldsymbol{x} \in \operatorname{cl}\mathcal{E}$ if for every $N_r(\boldsymbol{x})$, $N_r(\boldsymbol{x}) \cap \mathcal{E} \ne \varnothing$. $\mathcal{E}$ closed $\Leftrightarrow \mathcal{E} = \operatorname{cl}\mathcal{E}$

- The union of open sets is open. The intersection of a *finite* number of open sets is open. The intersection of closed sets is closed. The union of a *finite* number of closed sets is closed.
- The set of reals is both closed and open.
- ***Theorem***: $f^{-1}(\mathcal{A}) = \left\{\boldsymbol{x} \in \text{dom } f : f(\boldsymbol{x}) \in \mathcal{A}\right\}$. If dom f is open/closed and $\mathcal{A}$ is open/closed, then $f^{-1}(\mathcal{A})$ is also open/closed.

- $\mathcal{E} \subset \mathbb{R}^n$ is (sequentially) compact if for every sequence $\left\{\boldsymbol{x}_k\right\} \subseteq \mathcal{E}$ there exists a subsequence $\left\{\boldsymbol{x}_{k_i}\right\}$ converging to an element $\boldsymbol{x} \in \mathcal{E}$. [Another definition, equivalent in metric spaces, is that every open over must have a finite sub-cover].
  - ***Theorem*** (Heine-Borel): in finite dimensional spaces, a set is compact if and only if it is closed and bounded.
  - ***Theorems***: A closed subset of a compact set is compact. The intersection of a sequence of non-empty, nested compact sets is non-empty.
- The indicator function of a set $I_{\mathcal{C}}(\boldsymbol{x})$ is equal to 0 if $\boldsymbol{x} \in \mathcal{C}$ and ∞ otherwise.
- A *subspace* of a vector space is a subset of the vector space that contains the $\boldsymbol{0}$ vector and that satisfies closure under addition and scalar multiplication.

- ***Analysis & Calculus***
  - A function f is *coercive* over $\mathcal{C}$ if for every sequence $\{\boldsymbol{x}_k\} \subset \mathcal{C}$ with $\left\|\boldsymbol{x}_k\right\| \to \infty$, we have $\lim_{k\to\infty} f(\boldsymbol{x}_k) = \infty$.
  - Second-order Taylor expansion: if f is twice-continuously differentiable over $N_r(\boldsymbol{x})$, then $\forall\ \boldsymbol{d} \in N_r(\boldsymbol{0})$, $f(\boldsymbol{x}+\boldsymbol{d}) = f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top \boldsymbol{d} + \frac{1}{2}\boldsymbol{d}^\top \nabla^2 f(\boldsymbol{x})\boldsymbol{d} + o(\left\|\boldsymbol{d}\right\|^2)$
  - For a vector-value function $\boldsymbol{F}$, $\left[\nabla \boldsymbol{F}(\boldsymbol{x})\right]_{ij} = \partial F_j(\boldsymbol{x}) / \partial x_i$.
  - The chain rule states that $\nabla\left[\boldsymbol{g}(\boldsymbol{f}(\boldsymbol{x}))\right] = \nabla \boldsymbol{f}(\boldsymbol{x}) \nabla \boldsymbol{g}(\boldsymbol{f}(\boldsymbol{x}))$.

- ***Linear algebra***
  - The *range* or *image* of $A \in \mathbb{R}^{m\times n}$, $\mathcal{R}(A)$ is the set of all vectors that can be written as Ax. The *nullspace* or *kernel* $\mathcal{N}(A)$ is the set of vectors than satisfy $A\boldsymbol{x}$ = 0. Note that $\mathcal{N}(A) \overset{\perp}{\oplus} \mathcal{R}(A^\top) = \mathbb{R}^n$; in other words $\mathcal{R}(A) = \left[\mathcal{N}(A^\top)\right]^\top$. This last statement means that $\boldsymbol{z} = A\boldsymbol{x} \Leftrightarrow \boldsymbol{z}\cdot\boldsymbol{y} = 0\ \forall \boldsymbol{y}$ with $A^\top \boldsymbol{y} = \boldsymbol{0}$.

- A real symmetric matrix $A \in \mathbb{S}^n \subseteq \mathbb{R}^{n\times n}$ can be factored as $A = Q\Lambda Q^\top$, where $\Lambda = \operatorname{diag}(\lambda_1, \cdots, \lambda_n)$. Note that $\det A = \prod \lambda_i$, $\operatorname{tr} A = \sum \lambda_i$, $\|A\|_2 = \max|\lambda_i|$ and $\|A\|_F = \sqrt{\sum \lambda_i^2}$.
- $\mathbb{S}^n_{++}$ is the set of symmetric, positive definite matrix; all their eigenvectors are positive, and $A \in \mathbb{S}^n_{++}, \boldsymbol{x} \neq 0 \Rightarrow \boldsymbol{x}^\top A \boldsymbol{x} > 0$.
- Note that $\lambda_{\max/\min}(A) = \sup/\inf_{x\neq 0} \frac{x^\top A x}{x^\top x}$. As such, $\lambda_{\min}(A)x^\top x \leq x^\top A x \leq \lambda_{\max}(A)x^\top x$
- Suppose $A \in \mathbb{R}^{m\times n}$ with rank r. Then we can factor $A = U\Sigma V^\top$, where $U \in \mathbb{R}^{m\times r}, V \in \mathbb{R}^{n\times r}$ are two orthogonal matrices and $\Sigma = \operatorname{diag}(\sigma_1, \cdots, \sigma_r)$. Writing $A^\top A = V\Sigma U^\top U \Sigma V^\top = V\Sigma^2 V^\top$ we see that these singular values are the square root of the non-zero eigenvalues of $A^\top A$, and the *right singular vectors* V are the eigenvectors of $A^\top A$. Similarly, U contains the eigenvectors of $AA^\top$. We have $\sigma_{\max}(A) = \sup_{x,y\neq 0} \frac{x^\top A y}{\|x\|_2 \|y\|_2} = \sup_{y\neq 0} \frac{y^\top A^\top A y}{\|y\|_2}$. In other words, $\|A\|_2 = \sigma_{\max}(A)$. We denote $\operatorname{cond}(A) = \|A\|_2 \|A^{-1}\|_2 = \sigma_{\max}(A) / \sigma_{\min}(A)$.
- The pseudo-inverse is $A^\dagger = V\Sigma^{-1}U^\top$. If A is square and non-singular, then $A^\dagger = A^{-1}$. It is useful for the following reasons
  - $\boldsymbol{x} = A^\dagger \boldsymbol{b}$ is the minimum-Euclidean norm solution of $\min \|A\boldsymbol{x} - \boldsymbol{b}\|_2^2$.
  - The optimal value of $\min_{\boldsymbol{x}} \frac{1}{2}\boldsymbol{x}^\top P \boldsymbol{x} + \boldsymbol{q}\cdot\boldsymbol{x} + \boldsymbol{r}$ for $P \in \mathbb{S}^n$ can be expressed as $-\frac{1}{2}\boldsymbol{q}^\top P^\dagger \boldsymbol{q} + \boldsymbol{r}$ if $P \succeq 0$ and $\boldsymbol{q} \in \mathcal{R}(P)$, and $-\infty$ otherwise.

  Consider a matrix $X = \begin{bmatrix} A & B \\ B^\top & C \end{bmatrix} \in \mathbb{S}^n$, with $A \in \mathbb{S}^k$. If $\det A \neq 0$, then the matrix Then $S = C - B^\top A^{-1} B$ is the *Schur complement* of A in X. $\det X = \det A \det S$.
  - Let $A \succ 0$ and consider $\min_{\boldsymbol{u}} \boldsymbol{x}^\top X \boldsymbol{x} = \min_{u} \boldsymbol{u}^\top A \boldsymbol{u} + 2\boldsymbol{v}^\top B^\top \boldsymbol{u} + \boldsymbol{v}^\top C \boldsymbol{v}$, where $\boldsymbol{x} = \begin{bmatrix} \boldsymbol{u} & \boldsymbol{v} \end{bmatrix}$. This is a quadratic with solution $\boldsymbol{u} = -A^{-1}B\boldsymbol{v}$ and optimal value $\boldsymbol{v}^\top S \boldsymbol{v}$. Thus
    - $X \succ 0 \;\Leftrightarrow\; A \succ 0$ and $S \succ 0$
    - If $A \succ 0$, then $X \succeq 0 \Leftrightarrow S \succeq 0$

- Consider $\begin{bmatrix} A & B \\ B^\top & C \end{bmatrix}\begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{y} \end{bmatrix} = \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix}$ with $\det A \neq 0$. Using the top equation to eliminate x and feeding it into the bottom block, we get $y = S^{-1}\left(v - B^\top A^{-1} u\right)$. Substituting this back to find x, we get

$$\begin{bmatrix} A & B \\ B^\top & C \end{bmatrix}^{-1} = \begin{bmatrix} A^{-1} + A^{-1}BS^{-1}B^\top A^{-1} & -A^{-1}BS^{-1} \\ -S^{-1}B^\top A^{-1} & S^{-1} \end{bmatrix}$$