Statistical Inference via Convex Optimization

Anatoli B. Juditsky, Arkadi Nemirovski

► To cite this version:

Anatoli B. Juditsky, Arkadi Nemirovski. Statistical Inference via Convex Optimization. 3rd cycle. France. 2018. cel-02047265

HAL Id: cel-02047265

<https://hal.science/cel-02047265>

Submitted on 24 Feb 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Statistical Inference via Convex Optimization

Anatoli Juditsky and Arkadi Nemirovski

https://www.isye.gatech.edu/~nemirovs/StatOpt_LN.pdf

ENSAE ParisTech, October-November 2018

Transparencies in Dropbox <https://bit.ly/2AhVMNH>

Preface

- Many inference procedures in Statistics reduce to optimization

Example: MLE – Maximum Likelihood Estimation

Problem: *Given a parametric family $\{p_\theta(\cdot) : \theta \in \Theta\}$ of probability densities on $\mathbb{R}^d$ and a random observation ω drawn from some density $p_{\theta_\star}(\cdot)$ from the family, estimate the parameter $\theta_\star$.*

MLE: Given ω , maximize $p_\theta(\omega)$ over $\theta \in \Theta$ and use the maximizer $\hat{\theta} = \hat{\theta}(\omega)$ as an estimate of $\theta_\star$.

- In MLE, (most of the time) optimization is used for number crunching only and has nothing to do with motivation and performance analysis of MLE.
- *Most of traditional applications of Optimization in Statistics are of “number crunching” nature. Though important, “number crunching” applications are beyond our scope*

Our scope: inference routines motivated and justified by Optimization Theory

Our “working horse” will be Convex Optimization Theory – Convex Analysis, Optimality Conditions, Duality...

This choice is motivated by

- nice geometry of **convex** sets, functions, and optimization problems
- **computational tractability** of convex optimization implying **computational efficiency** of statistical inferences stemming from Convex Optimization.

Course validation

Date: *return your copy by Thursday, November 29*

Options: *Any choice of two (simple) or one (more involved, marked with †) exercises from Chapters 2, 3, or 4 of the manuscript*

https://www2.isye.gatech.edu/~nemirovs/StatOpt_LN.pdf

*FOR INTRODUCTION:
LINEAR FUNCTIONAL ESTIMATION
IN WHITE NOISE*

- *Donoho's Theorem*
- *Estimating Linear Functionals on Unions of Convex Sets*

Linear Functional Estimation in White Noise Model

We consider the model

$$\omega = Ax + \sigma\xi$$

with

- *observable* $\omega \in \mathbb{R}^d$;
- *unknown parameter* $x \in \mathbb{R}^n$; it is known *a priori* that $x \in X$, $X \subset \mathbb{R}^n$ being *convex and compact*;
- *sensing matrix* $A \in \mathbb{R}^{d \times n}$;
- *observation noise* $\xi \sim N(0, I_d)$, $\sigma > 0$ is known.

Our goal is to estimate the value $g(x)$ at x of a given linear functional

$$g(z) = g^T z.$$

Estimation risk: given $\epsilon \in (0, 1)$, we define the risk of an estimator $\hat{g}(\omega)$ of $g^T x$ on X as *the maximal width of confidence interval of level $1 - \epsilon$* :

$$\text{Risk}_\epsilon[\hat{g}|X] = \min \left\{ \rho : \text{Prob}_x \left\{ |\hat{g}(\omega) - g^T x| > \rho \right\} \leq \epsilon \forall x \in X \right\}$$

We are specifically interested in the *minmax risk*

$$\text{RiskOpt}_\epsilon = \inf_{\hat{g}} \text{Risk}_\epsilon[\hat{g}|X],$$

- We say that an *estimator* $\hat{g} = \hat{g}_{c,\phi}$ is *affine* if it is of the form

$$\hat{g}_{c,\phi}(\omega) = \phi^T \omega + c, \quad \phi \in \mathbb{R}^d, \quad c \in \mathbb{R}.$$

We want to compare the minmax risk RiskOpt_ϵ to the *maximal on X risk of the best affine estimator*:

$$\text{RiskA}_\epsilon = \inf_{c,\phi} \text{Risk}_\epsilon[\hat{g}_{c,\phi}|X].$$

White Noise Model : Donoho's Theorem

Theorem [Donoho, 1995] *In the white noise model*

$$\text{Risk}A_\epsilon \leq C \text{RiskOpt}_\epsilon$$

with $C \leq 1.25$ when $\epsilon \leq 0.01$.

Let us prove the following version of this result:

Proposition *Let $\epsilon \in (0, 1/2)$. There is an affine estimator $\hat{g}_{c,\phi}(\cdot)$ of $g(x)$ such that*

$$\text{Risk}_\epsilon[\hat{g}_{c,\phi}|X] \leq \frac{q_{\mathcal{N}}(1 - \epsilon/2)}{q_{\mathcal{N}}(1 - \epsilon)} \text{RiskOpt}_\epsilon$$

where $q_{\mathcal{N}}(\alpha)$ is the α -quantile of $\mathcal{N}(0, 1)$.

Note: *the ratio $\frac{q_{\mathcal{N}}(1-\epsilon/2)}{q_{\mathcal{N}}(1-\epsilon)} \rightarrow 1$ as $\epsilon \downarrow +0$*

Proof. 1^o Note that $\forall u \in X$,

$$\begin{aligned}\widehat{g}_{c,\phi}(\omega) - g(u) &= \phi^T \omega + c - g^T u = \phi^T A u - g^T u + \underbrace{\sigma \phi^T \xi}_{\sim \mathcal{N}(0, \sigma^2 \|\phi\|_2^2)} \\ &\leq (A^T \phi - g)u + c + q_{\mathcal{N}}(1 - \epsilon/2) \sigma \|\phi\|_2\end{aligned}$$

with probability $1 - \epsilon/2$.

Thus, for $\rho = \max_{u \in X} (A^T \phi - g)u + c + q_{\mathcal{N}}(1 - \epsilon/2) \sigma \|\phi\|_2$,

$$\sup_{u \in X} \text{Prob}_u \{ \widehat{g}(\omega) - g(u) \geq \rho \} \leq \epsilon/2.$$

For the same reasons,

$$\sup_{v \in X} \text{Prob}_v \{ \widehat{g}(\omega) - g(v) \leq -\rho' \} \leq \epsilon/2.$$

where $\rho' = \max_{v \in X} (g - A^T \phi)v - c + q_{\mathcal{N}}(1 - \epsilon/2) \sigma \|\phi\|_2$.

Let us put

$$c = \frac{1}{2} \left[\max_{u \in X} (g - A^T \phi)^T u - \max_{v \in X} (A^T \phi - g)^T v \right].$$

With this choice of c we have $\rho = \rho'$, and

$$\rho = \frac{1}{2} \max_{u, v \in X} \{ \Phi(\phi; u, v) := [g^T(v - u) + \phi^T A(u - v)] + 2q_{\mathcal{N}}(1 - \epsilon/2) \sigma \|\phi\|_2 \}.$$

is the half-length of the confidence interval of $\widehat{g}(Y)$. We conclude that

$$\text{Risk}_{\epsilon}[\widehat{g}_{c,\phi}|X] \leq \text{Opt}_X := \min_{\phi} \left[\max_{u, v \in X} \Phi(\phi; u, v) \right]$$

2° Observe that $\Phi(\phi; u, v)$ is convex and coercive in ϕ and concave in u, v . Thus, by the *Sion-Kakutani theorem*,

$$\min_{\phi} \left[\max_{u, v \in X} \Phi(\phi; u, v) \right] = \max_{u, v \in X} \left[\min_{\phi} \Phi(\phi; u, v) \right]$$

Note that

$$\begin{aligned} \min_{\phi} \Phi(\phi; u, v) &= g^T(v - u) + \min_{\phi} [\phi^T A(u - v) + 2q_N(1 - \epsilon/2)\sigma \|\phi\|_2] \\ &= \begin{cases} g^T(v - u), & \|A(u - v)\|_2 \leq 2q_N(1 - \epsilon/2)\sigma, \\ -\infty, & \text{otherwise,} \end{cases} \end{aligned}$$

and we conclude that

$$\text{Opt}_X = \max \left\{ \frac{1}{2}g^T(v - u) : u, v \in X, \|A(v - u)\|_2 \leq 2q_N(1 - \epsilon/2)\sigma \right\}. \quad (*)$$

3° We need to show that

$$\text{Opt}_X \leq \frac{q_N(1 - \epsilon/2)}{q_N(1 - \epsilon)} \text{RiskOpt}_{\epsilon} := \psi(\epsilon) \text{RiskOpt}_{\epsilon}.$$

Let us assume, on the contrary, that

$$\rho_* = \text{RiskOpt}_{\epsilon} < \psi(\epsilon)^{-1} \text{Opt}_X.$$

Observe that $(*)$ is clearly solvable, let $[\bar{u}; \bar{v}]$ be an optimal solution; we have

$$2\rho_* < \psi(\epsilon)^{-1} g^T(\bar{v} - \bar{u}) < g^T(\bar{v} - \bar{u}).$$

Let now $[\hat{u}; \hat{v}]$ be a “contraction” of $[\bar{u}; \bar{v}]$ such that

$$2\rho_* = g^T(\hat{v} - \hat{u}).$$

We have

$$\|A(\hat{v} - \hat{u})\|_2 = \underbrace{\frac{\rho_*}{\text{Opt}_X}}_{< \psi^{-1}} \|A(\bar{v} - \bar{u})\|_2 < 2q_N(1 - \epsilon)\sigma.$$

Now let us consider the problem of deciding between hypotheses

$$H_1 : x = \hat{u} \text{ and } H_2 : x = \hat{v}.$$

Let $g_*(\cdot)$ be the (minmax) estimator of $g(u)$ with risk ρ_* .

Now consider the rule $\mathcal{T}(\cdot)$ deciding between H_0 and H_1 :

$\mathcal{T}$ decides H_1 if $g_(\omega) \leq \frac{1}{2}g^T(\hat{v} + \hat{u})$, and decides H_2 otherwise.*

By construction, the probability that $\mathcal{T}(\omega)$ rejects H_i , $i = 1, 2$ when it is true is $\leq \epsilon$.
On the other hand,

$$\|A\hat{v} - A\hat{u}\|_2 < 2q_N(1 - \epsilon)\sigma.$$

Then, by the *Neyman-Pearson lemma*, there is no test for H_1 and H_1 with the risk 2ϵ , and we arrive at a contradiction. $\square$

Interpreting the result

In order to find $\hat{g}$ one has to solve a conic convex problem

$$\text{Opt}_X = \max_{u,v \in X} \left\{ \frac{1}{2}g^T(v-u) : \underbrace{\|A(v-u)\|_2 - 2q_{\mathcal{N}}(1 - \epsilon/2)\sigma}_{:=f(u,v)} \leq 0 \right\}.$$

It is a “well structured” problem and can be solved efficiently if X is *computationally tractable*.

Note that an optimal solution $(\bar{u}, \bar{v})$ to this problem can be augmented by the Lagrange multiplier $\nu \geq 0$ such that the vectors

$$\begin{aligned} e_u &= \left. \frac{\partial}{\partial u} \right|_{(u,v)=(\bar{u},\bar{v})} \left[\frac{1}{2}g^T u + \nu f(u,v) \right] \\ e_v &= \left. \frac{\partial}{\partial v} \right|_{(u,v)=(\bar{u},\bar{v})} \left[-\frac{1}{2}g^T v + \nu f(u,v) \right] \end{aligned}$$

belong to the normal cone of $X \times X$ at $(\bar{u}, \bar{v})$:

$$\forall (u, v \in X) : e_u^T(u - \bar{u}) \geq 0, e_v^T(v - \bar{v}) \geq 0. \quad (!)$$

There are two possible cases:

- $\nu = 0$: in this case the constraint $\|A(v - u)\|_2 \leq 2q_{\mathcal{N}}(1 - \epsilon/2)$ is not *active*, and

$$[\bar{u}; \bar{v}] \in \text{Argmax} \left\{ g^T(v - u) : u, v \in X \right\}.$$

We have $\phi_* = 0$ and $c_* = g^T(\bar{u} + \bar{v})$;

- $\nu > 0$: in this case $\|A(\bar{v} - \bar{u})\|_2 = 2q_{\mathcal{N}}(1 - \epsilon/2)\sigma$, and we set

$$\hat{g}(\omega) = \phi_*^T \omega + c_*, \quad \phi_* = \nu \frac{A(\bar{v} - \bar{u})}{\|A(\bar{u} - \bar{v})\|_2}, \quad c_* = \frac{1}{2} g^T(\bar{u} + \bar{v}).$$

One verifies straightforwardly that $\Phi(\phi_; u, v)$ attains its maximum over $u, v \in X$ at the point $(\bar{u}, \bar{v})$, and the maximal value is exactly 2Opt_X .*

Linear Functional Estimation over Unions of Convex Sets

- *Given*
 - collection $X_j, j = 1, \dots, J$, of nonempty convex compact sets in $\mathbb{R}^n$
 - $d \times n$ matrix A
 - linear function $g^T x$ on $\mathbb{R}^n$
 - observation $\omega = Ax + \xi, \xi \sim \mathcal{N}(0, \sigma^2 I_d), \sigma > 0$,

we want to estimate $g^T x$.

Same as before, we quantify the performance of an estimate $\hat{g}(\cdot)$ by its ϵ -*risk*:

$$\text{Risk}_\epsilon[\hat{g}|X] = \min \left\{ \rho : \text{Prob}_x \left\{ |\hat{g}(\omega) - g^T x| > \rho \right\} \leq \epsilon \forall (j \leq J, x \in X_j) \right\}.$$

We consider estimates which are aggregates of “pairwise” estimates of the form

$$g_{ij}(\omega) = \phi_{ij}^T \omega + c_{ij}, \quad 1 \leq i, j \leq J.$$

Let $\varepsilon \in (0, 1)$, $g(x) = \phi^T \omega + c$, $\phi \in \mathbb{R}^d$, $c \in \mathbb{R}$. Observe that $\forall x \in X$,

$$\begin{aligned} g(\omega) - g^T x &= (A^T \phi - g)x + c + \sigma \phi^T \xi \leq (A^T \phi - g)x + c + \sigma \|\phi\|_2 q_N (1 - \varepsilon/2), \\ g^T x - g(\omega) &= (g - A^T \phi)x - c + \sigma \phi^T \xi \leq (g - A^T \phi)x - c + \sigma \|\phi\|_2 q_N (1 - \varepsilon/2) \end{aligned}$$

with probability $1 - \varepsilon/2$. Therefore, for any $1 \leq i, j \leq J$,

$$\begin{aligned} \forall u \in \mathbf{X}_i: \quad g(\omega) - g^T u &\leq \max_{u \in \mathbf{X}_i} (A^T \phi - g)u + c + \sigma \|\phi\|_2 q_N (1 - \varepsilon/2), \\ \forall v \in \mathbf{X}_j: \quad g^T v - g(\omega) &\leq \max_{v \in \mathbf{X}_j} (g - A^T \phi)v - c + \sigma \|\phi\|_2 q_N (1 - \varepsilon/2). \end{aligned}$$

Let

$$c = \frac{1}{2} \left[\max_{v \in \mathbf{X}_j} (g - A^T \phi)v - \max_{u \in \mathbf{X}_i} (A^T \phi - g)u \right].$$

We conclude that $g(\omega)$ satisfies:

$$\begin{aligned} \sup_{u \in \mathbf{X}_i} \text{Prob}_u \{g(\omega) - g^T u \geq \rho\} &\leq \varepsilon/2, \\ \sup_{v \in \mathbf{X}_j} \text{Prob}_v \{g^T v - g(\omega) \geq \rho\} &\leq \varepsilon/2, \end{aligned}$$

where

$$\rho = \frac{1}{2} \max_{u \in \mathbf{X}_i, v \in \mathbf{X}_j} \left\{ \Phi(\phi; u, v) := [g^T(v - u) + \phi^T A(u - v)] + 2q_N(1 - \varepsilon/2)\sigma \|\phi\|_2 \right\}.$$

Building the estimate: Let $\varepsilon = \epsilon/J$. For $1 \leq i, j \leq J$, consider convex optimization problems

$$\begin{aligned} \text{Opt}_{ij} &= \min_{\phi} \left\{ \max_{u \in X_i, v \in X_j} \frac{1}{2} \Phi(\phi; u, v) \right\} \\ &= \min_{\phi} \left\{ \max_{u \in X_i, v \in X_j} \frac{1}{2} \left[(g^T(v - u) + \phi^T A(u - v)) + 2q_{\mathcal{N}} \left(1 - \frac{\epsilon}{2J}\right) \sigma \|\phi\|_2 \right] \right\}. \end{aligned}$$

Let ϕ_{ij} be feasible solutions to these problems, and ρ_{ij} be the corresponding values of the objectives (the closer to Opt_{ij} , the better).

- We set

$$g_{ij}(\omega) = \phi_{ij}^T \omega + c_{ij}, \quad c_{ij} = \frac{1}{2} \left[\max_{v \in X_j} (g - A^T \phi)v - \max_{u \in X_i} (A^T \phi - g)u \right].$$

thus ensuring

$$\begin{aligned} \text{Prob}_u \{g_{ij}(\omega) > g^T u + \rho_{ij}\} &\leq \frac{\epsilon}{2J}, \quad u \in X_i \\ \text{Prob}_v \{g_{ij}(\omega) < g^T v - \rho_{ij}\} &\leq \frac{\epsilon}{2J}, \quad v \in X_j \end{aligned}$$

$$\text{Prob}_u \{g_{ij}(\omega) > g^T u + \rho_{ij}\} \leq \frac{\epsilon}{2J}, \quad u \in X_i$$

$$\text{Prob}_v \{g_{ij}(\omega) < g^T v - \rho_{ij}\} \leq \frac{\epsilon}{2J}, \quad v \in X_j$$

- Given observation ω , we build the matrix

$$\left[E_{ij} := g_{ij}(\omega) \right]_{1 \leq i, j \leq J},$$

and set

$$\begin{aligned} r_i &= \max_j E_{ij}, \quad r = \min_i r_i = \min_i \max_j E_{ij}, \\ s_j &= \min_i E_{ij}, \quad s = \max_j s_j = \max_j \min_i E_{ij}. \end{aligned}$$

The estimate $\hat{g}(\omega)$ of $g^T x$, $x \in X = \cup_j X_j$, is set to

$$\hat{g}(\omega) = \frac{1}{2}[r + s].$$

Theorem ϵ -Risk of the estimate $\hat{g}$ satisfies

$$\text{Risk}_\epsilon[\hat{g}|X] \leq \rho := \max_{i,j} \rho_{ij}.$$

$$\text{Prob}_u \left\{ g_{ij}(\omega) > g^T u + \rho_{ij} \right\} \leq \frac{\epsilon}{2J}, \quad u \in X_i \quad (a)$$

$$\text{Prob}_v \left\{ g_{ij}(\omega) < g^T v - \rho_{ij} \right\} \leq \frac{\epsilon}{2J}, \quad v \in X_j \quad (b)$$

$$E_{ij} = g_{ij}(\omega) \quad (c)$$

$$r_i = \max_j E_{ij}, \quad r = \min_i r_i = \min_i \max_j E_{ij}, \quad (c)$$

$$s_j = \min_i E_{ij}, \quad s = \max_j s_j = \max_j \min_i E_{ij} \quad (d)$$

$$\widehat{g}(\omega^K) = \frac{1}{2}[r + s] \quad (d)$$

Proof. 1⁰. By (b),

$$r_i \geq E_{ij} \geq s_j \Rightarrow r_i \geq s_j \quad \forall i, j \quad (e)$$

2⁰. Let $x \in X_\ell$, and let $\mathcal{E}$ be the event

$$g_{\ell j}(\omega) \leq g^T x + \rho_{\ell j} \quad \forall j \leq J, \quad \text{and} \quad g_{i\ell}(\omega) \geq g^T x - \rho_{i\ell} \quad \forall i \leq J,$$

so that by (a) and union bound,

$$P(\mathcal{E}) \geq 1 - \epsilon.$$

Let $\Delta = [g^T x - \rho, g^T x + \rho]$. *Note that when $\omega \in \mathcal{E}$, one has $r, c \in \Delta$.*

Indeed, when $\omega \in \mathcal{E}$, we have

$$r_{\ell j} \leq g^T x + \rho_{\ell j} \quad \forall j \Rightarrow r_\ell \leq g^T x + \rho \Rightarrow r \leq g^T x + \rho,$$

same as

$$s_{i\ell} \geq g^T x - \rho_{i\ell} \quad \forall i \Rightarrow s_\ell \geq g^T x - \rho \Rightarrow s \geq g^T x - \rho.$$

Since $r \geq s$ by (e), we conclude that $r, s \in \Delta$ whenever $\omega \in \mathcal{E}$. □

Near-Optimality

$$\text{Opt}_{ij} = \min_{\phi} \left\{ \max_{u \in X_i, v \in X_j} \underbrace{\frac{1}{2} \left[(g^T(v - u) + \phi^T A(u - v)) + 2q_{\mathcal{N}} \left(1 - \frac{\epsilon}{2J} \right) \sigma \|\phi\|_2 \right]}_{\Phi(\phi_{ij}; u, v)} \right\}$$

Because $\Phi(\phi_{ij}; u, v)$ is convex and coercive in ϕ and concave in u, v while X_i, X_j are convex compact sets, by the *Sion-Kakutani theorem*,

$$\min_{\phi} \left[\max_{u \in X_i, v \in X_j} \Phi(\phi_{ij}; u, v) \right] = \max_{u \in X_i, v \in X_j} \left[\min_{\phi} \Phi(\phi; u, v) \right]$$

Same as above, we conclude that Opt_{ij} is either $-\infty$, or

$$\text{Opt}_{ij} = \max \left\{ \frac{1}{2} g^T(v - u) : u \in X_i, v \in X_j, \|A(v - u)\|_2 \leq 2q_{\mathcal{N}} \left(1 - \frac{\epsilon}{2J} \right) \sigma \right\}. \quad (*)$$

It is immediately seen that Opt_{ii} is a nonnegative real, so that

$$\text{Opt} := \max_{i,j} \text{Opt}_{ij} \geq 0.$$

Note that we can make the bound ρ on ϵ -risk of our estimate arbitrarily close to Opt by properly selecting corresponding feasible solutions.

We are to show that Opt *is nearly the optimal ϵ -risk* in our estimation problem

$$\text{RiskOpt}_\epsilon = \inf_{\hat{g}} \text{Risk}_\epsilon[\hat{g}|X].$$

Theorem *Let $\epsilon \in (0, 1/2)$. Assume that we are in the special case where for every i, j there exists $x_{ij} \in X_i \cap X_j$. Then*

$$\text{Opt} \leq \frac{q_{\mathcal{N}}(1 - \frac{\epsilon}{2J})}{q_{\mathcal{N}}(1 - \epsilon)} \text{RiskOpt}_\epsilon.$$

$$\text{Opt}_{ij} = \max \left\{ \frac{1}{2} g^T(v - u) : u \in X_i, v \in X_j, \|A(v - u)\|_2 \leq 2q_{\mathcal{N}} \left(1 - \frac{\epsilon}{2J}\right) \sigma \right\}. \quad (*)$$

Proof. 1^o We are to show that for all $1 \leq i, j \leq J$

$$\text{Opt}_{ij} \leq \underbrace{\frac{q_{\mathcal{N}}(1 - \frac{\epsilon}{2J})}{q_{\mathcal{N}}(1 - \epsilon)}}_{:=\psi(\epsilon)} \text{RiskOpt}_{\epsilon} [\geq 0]$$

Let us associate with (*) a *parametric problem*

$$w_{ij}(s) = \max \left\{ \frac{1}{2} g^T(v - u) : u \in X_i, v \in X_j, \|A(v - u)\|_2 \leq s \right\}. \quad (P_s)$$

Note that (P_s) is *convex* (in the sense that its objective is concave and the domain is convex) and *feasible* $\forall s \geq 0$ (with a feasible solution $u = v = x_{ij}$) with $w_{ij}(0) \geq 0$. We conclude that $w_{ij}(s)$ is nonnegative, concave and bounded on $\mathbb{R}_+ \Rightarrow w_{ij}(s)$ is continuous for $s > 0$.

$$w_{ij}(s) = \max \left\{ \frac{1}{2} g^T(v - u) : u \in X_i, v \in X_j, \|A(v - u)\|_2 \leq s \right\}. \quad (P_s)$$

We claim that

$$w_{ij}(s_*) \leq \text{RiskOpt}_\epsilon, \quad s_* = 2\sigma q_{\mathcal{N}}(1 - \epsilon).$$

Let us assume, on the contrary, that $w_{ij}(s_*) > \text{RiskOpt}_\epsilon$, so that for some $s' \in (0, s_*)$, $w_{ij}(s') > \text{RiskOpt}_\epsilon$. Then there exist $\bar{u} \in X_i$ and $\bar{v} \in X_j$ which are feasible for $(P_{s'})$ and such that

$$\|A(\bar{u} - \bar{v})\|_2 < 2\sigma q_{\mathcal{N}}(1 - \epsilon),$$

and

$$\frac{1}{2} g^T(\bar{v} - \bar{u}) > \text{RiskOpt}_\epsilon.$$

We have already seen that this is impossible by the Neyman-Pearson lemma, thus arriving to a contradiction.

2°. We conclude that $w_{ij}(s_*) \leq \text{RiskOpt}_\epsilon$. Now, we for

$$\bar{s} = \psi(\epsilon) s_* = 2q_{\mathcal{N}} \left(1 - \frac{\epsilon}{2J} \right) \sigma$$

we have by concavity of $w_{ij}(s)$ and due to $w_{ij}(0) \geq 0$:

$$w_{ij}(\bar{s}) \leq \psi(\epsilon) w_{ij}(s') \leq \psi(\epsilon) \text{RiskOpt}_\epsilon. \quad \square$$

Numerical Illustration

Setup: Given $J = 100$ points $x_j \in \mathbb{R}^{20}$ and stationary K -repeated observation noisy observation

$$\omega^K = (\omega_1, \dots, \omega_K), \quad \omega_k \sim \mathcal{N}(Ax, I_{20})$$

of one of the points (we do not know which one!), we want to recover the first entry of the point.

- A : randomly generated matrix
- $\epsilon = 0.01$.

Note: We are in the situation where $X_i = \{x_i\}$ are singletons.

$\text{Opt}(K)$ vs. σ^{-2} , data over 20 randomly generated collections $\{x_i\}_{i=1}^{100}$

HYPOTHESIS TESTING, I

- *Preliminaries*
 - *Tests & Risks*
 - *Repeated Observations*
 - *2-Point Lower Risk Bound*
- *Pairwise Tests via Euclidean Separation*
- *From Pairwise to Multiple Hypothesis Testing*

Hypothesis Testing Problem

Given

- *observation space* Ω where our observations take values,
 - L *families* $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_L$ of probability distributions on Ω , and
 - *an observation* ω – a realization of random variable with *unknown* probability distribution P *known to belong to one of the families* $\mathcal{P}_\ell$: $P \in \bigcup_{\ell=1}^L \mathcal{P}_\ell$,
- we want to decide to which one of the families $\mathcal{P}_\ell$ the distribution P belongs.

Equivalent wording: *Given the outlined data, we want to decide on L hypotheses $H_1, \dots, H_L$, with ℓ -th hypothesis H_ℓ stating that $P \in \mathcal{P}_\ell$.*

A test is a function $\mathcal{T}(\cdot)$ on Ω . The value $\mathcal{T}(\omega)$ of this function at a point $\omega \in \Omega$ is a subset of the set $\{1, \dots, L\}$.

- relation $\ell \in \mathcal{T}(\omega)$ is interpreted as “given observation ω , the test accepts the hypothesis H_ℓ ”
- relation $\ell \notin \mathcal{T}(\omega)$ is interpreted as “given observation ω , the test rejects the hypothesis H_ℓ ”

Simple tests

- $\mathcal{T}$ is called *simple*, if $\mathcal{T}(\omega)$ is a singleton for every $\omega \in \Omega$.

Risks of a simple test $\mathcal{T}$, are defined as follows:

- ℓ -th partial risk of $\mathcal{T}$ is the (worst-case) probability to reject ℓ -th hypothesis when it is true:

$$\text{Risk}_\ell[\mathcal{T}|H_1, \dots, H_L] = \sup_{P \in \mathcal{P}_\ell} \text{Prob}_{\omega \sim P} \{\ell \notin \mathcal{T}(\omega)\}$$

- total risk of $\mathcal{T}$ is the sum of all partial risks:

$$\text{Risk}_{\text{tot}}[\mathcal{T}|H_1, \dots, H_L] = \sum_{1 \leq \ell \leq L} \text{Risk}_\ell[\mathcal{T}|H_1, \dots, H_L].$$

- risk of $\mathcal{T}$ is the maximum of all partial risks:

$$\text{Risk}[\mathcal{T}|H_1, \dots, H_L] = \max_{1 \leq \ell \leq L} \text{Risk}_\ell[\mathcal{T}|H_1, \dots, H_L].$$

Remark: What was called test is in fact a *deterministic* test.

A *randomized* test is a *deterministic* function $\mathcal{T}(\omega, \eta)$ of observation ω and *independent of ω random variable* $\eta \sim P_\eta$ with once for ever fixed distribution (e.g., $P_\eta = \text{Uniform}[0, 1]$).

- Values $\mathcal{T}(\omega, \eta)$ of $\mathcal{T}$ are subsets of $\{1, \dots, L\}$ (singletons for a simple test).
- Given observation ω , we “flip a coin” (draw a realization of η), accept hypotheses H_ℓ , $\ell \in \mathcal{T}(\omega, \eta)$, and reject all other hypotheses.
- Partial risks of a randomized test are

$$\text{Risk}_\ell[\mathcal{T}|H_1, \dots, H_L] = \sup_{P \in \mathcal{P}_\ell} \text{Prob}_{(\omega, \eta) \sim P \times P_\eta} \{\ell \notin \mathcal{T}(\omega, \eta)\}.$$

Exactly as above, these risks give rise to the *total risk* and *risk* of $\mathcal{T}$.

Testing from repeated observations

There are situations where an inference can be based on several observations $\omega_1, \dots, \omega_K$ rather than on a single observation. Our related setup is as follows:

We are given L families $\mathcal{P}_\ell$, $\ell = 1, \dots, L$, of probability distributions on observation space Ω and a collection

$$\omega^K = (\omega_1, \dots, \omega_K),$$

and want to make conclusions on how the distribution of ω^K “is positioned” w.r.t. the families $\mathcal{P}_\ell$, $1 \leq \ell \leq L$.

Specifically, we are interested in three situations of this type.

A. Stationary K -repeated observations

Situation: $\omega_1, \dots, \omega_K$ are drawn *independently of each other* from the same distribution P . Our goal is to decide, given ω^K , on the hypotheses $P \in \mathcal{P}_\ell$, $\ell = 1, \dots, L$.

Equivalently: Families $\mathcal{P}_\ell$ of probability distributions of $\omega \in \Omega$, $1 \leq \ell \leq L$, give rise to the families

$$\mathcal{P}_\ell^K = \{P^K = \underbrace{P \times \dots \times P}_K : P \in \mathcal{P}_\ell\}$$

of probability distributions on $\Omega^K = \underbrace{\Omega \times \dots \times \Omega}_K$. Given observation $\omega^K \in \Omega^K$, we want to decide on the hypotheses

$$H_\ell^K : \omega^K \sim P^K \in \mathcal{P}_\ell^K, \quad 1 \leq \ell \leq L.$$

B. Semi-stationary K -repeated observations

Situation: “The nature” selects somehow a sequence $P_1, \dots, P_K$ of distributions on Ω , and then draws, *independently across k* , observations ω_k from these distributions:

$$\omega_k \sim P_k \text{ are independent across } k \leq K$$

Our goal is to decide, given $\omega^K = (\omega_1, \dots, \omega_K)$, on the hypotheses

$$\{P_k \in \mathcal{P}_\ell, 1 \leq k \leq K\}, \ell = 1, \dots, L.$$

Equivalently: Families $\mathcal{P}_\ell$ of probability distributions of $\omega \in \Omega$, $1 \leq \ell \leq L$, give rise to the families

$$\mathcal{P}_\ell^{\oplus, K} = \bigoplus_{k=1}^K \mathcal{P}_\ell := \{P^K = P_1 \times \dots \times P_K : P_k \in \mathcal{P}_\ell, 1 \leq k \leq K\}$$

of probability distributions on $\Omega^K = \underbrace{\Omega \times \dots \times \Omega}_K$. Given observation $\omega^K \in \Omega^K$, we want to decide on the hypotheses

$$H_\ell^{\oplus, K} : \omega^K \sim P^K \in \mathcal{P}_\ell^{\oplus, K}, 1 \leq \ell \leq L.$$

C. Quasi-stationary K -repeated observations:

We observe random sequence $\omega^K = (\omega_1, \dots, \omega_K)$ generated as follows:

*There exists a random sequence $\zeta_1, \dots, \zeta_K$ of **driving factors** such that for $1 \leq k \leq K$*

- ω_k is a deterministic function of $\zeta^k = (\zeta_1, \dots, \zeta_k)$
- *conditional, given ζ^{k-1} , distribution of ω_k always belongs to $\mathcal{P}_\ell$.*

Our goal is to decide, given ω^K , on the underlying ℓ .

Equivalently: Families $\mathcal{P}_\ell$ of probability distributions on Ω , $1 \leq \ell \leq L$, give rise to the **quasi-direct products** $\mathcal{P}_\ell^{\otimes, K} = \bigotimes_{k=1}^K \mathcal{P}_\ell$ of families $\mathcal{P}_\ell$.

Family $\bigotimes_{k=1}^K \mathcal{P}_\ell$ is comprised of all probability distributions on $\Omega^K = \underbrace{\Omega \times \dots \times \Omega}_K$

which can be obtained from $\mathcal{P}_\ell$ via the above “driving factors” mechanism.

Given observation $\omega^K \in \Omega^K$, we want to decide on the hypotheses

$$H_\ell^{\otimes, K} : \omega^K \sim P^K \in \mathcal{P}_\ell^{\otimes, K}, \quad 1 \leq \ell \leq L.$$

Important fact: 2-point lower risk bound

Consider a *simple pairwise test* deciding on two **simple** hypotheses on the distribution P of observation $\omega \in \Omega$:

$$H_1 : P = P_1, \quad H_2 : P = P_2.$$

Let P_1, P_2 have densities p_1, p_2 w.r.t. some reference measure Π on Ω .

Then *the total risk of every test $\mathcal{T}$ deciding on H_1, H_2 admits lower bound as follows:*

$$\text{Risk}_{\text{tot}}[\mathcal{T}|H_1, H_2] \geq \int_{\Omega} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega).$$

As a result,

$$\text{Risk}[\mathcal{T}|H_1, H_2] \geq \frac{1}{2} \int_{\Omega} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega). \quad (*)$$

The bound does not depend on the choice of Π (for example, we can always take $\Pi = P_1 + P_2$).

$$\text{Risk}[\mathcal{T}|H_1, H_2] \geq \frac{1}{2} \int_{\Omega} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega). \quad (?)$$

Proof (for deterministic test). Simple test deciding on H_1, H_2 must accept H_1 and reject H_2 on some subset Ω_1 of Ω and must reject H_1 and accept H_2 on the complement $\Omega_2 = \Omega \setminus \Omega_1$ of this set. We have

$$\begin{aligned} \text{Risk}_1[\mathcal{T}|H_1, H_2] &= \int_{\Omega_2} p_1(\omega) \Pi(d\omega) \geq \int_{\Omega_2} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega) \\ \text{Risk}_2[\mathcal{T}|H_1, H_2] &= \int_{\Omega_1} p_2(\omega) \Pi(d\omega) \geq \int_{\Omega_1} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega), \end{aligned}$$

thus

$$\begin{aligned} \text{Risk}_{\text{tot}}[\mathcal{T}|H_1, H_2] &\geq \int_{\Omega_2} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega) + \int_{\Omega_1} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega) \\ &= \int_{\Omega} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega) \quad \square \end{aligned}$$

Remark: In the situation in question, consider the test $\mathcal{T}_*$:

$\mathcal{T}_*$ accepts H_1 if $p_1(\omega) \geq p_2(\omega)$, and accepts H_2 if $p_1(\omega) < p_2(\omega)$

(Neyman-Pearson or **Likelihood Ratio test**).

Lemma [Neyman-Pearson] *For any test $\mathcal{T}$*

$$\text{Risk}_{\text{tot}}[\mathcal{T}|H_1, H_2] \geq \text{Risk}_{\text{tot}}[\mathcal{T}_*|H_1, H_2] = \int_{\Omega} \min[p_1(\omega), p_2(\omega)] \Pi(d\omega).$$

Example: Gaussian distributions. Let $\Omega = \mathbb{R}^d$, our observation

$$\omega = x + \xi,$$

where x is deterministic and $\xi \sim \mathcal{N}(0, I)$, i.e., has the density $p(u) = (2\pi)^{-d/2} e^{-\frac{1}{2}u^T u}$. Our goal is to decide on two simple hypotheses: $H_1 : x = x_1$ and $H_2 : x = x_2$, x_1, x_2 being two given points.

The corresponding test $\mathcal{T}_*$ decides H_1 if

$$\phi_*(\omega) = e^T [\omega - \frac{1}{2}(x_1 + x_2)] \geq 0, \quad e = \frac{x^1 - x^2}{\|x^1 - x^2\|_2}$$

and decides H_2 otherwise. Its total risk is given by

$$\begin{aligned} \text{Risk}_{\text{tot}}[\mathcal{T}|H_1, H_2] &= P_1(\phi_*(\omega) < 0) + P_2(\phi_*(\omega) \geq 0) \\ &= 2\text{Prob}_{\zeta \sim N(0,1)} \left\{ \zeta \geq \frac{1}{2} \|x_2 - x_1\|_2 \right\} = \text{erfc} \left(\frac{\|x_2 - x_1\|_2}{2\sqrt{2}} \right). \end{aligned}$$

$$e = \frac{x^1 - x^2}{\|x^1 - x^2\|_2}$$

$$\phi(\omega) = e^T \omega - \underbrace{\frac{1}{2} e^T [x^1 + x^2]}_c$$

$$\Pi = \{\omega : \|\omega - x^1\|_2 = \|\omega - x^2\|_2\} = \{\omega : \phi(\omega) = 0\}$$

Corollary. Consider L hypotheses $H_\ell : P \in \mathcal{P}_\ell$, $\ell = 1, 2, \dots, L$, on the distribution P of observation $\omega \in \Omega$, let $\ell \neq \ell'$ and let $P_\ell \in \mathcal{P}_\ell$, $P_{\ell'} \in \mathcal{P}_{\ell'}$.

The risk of any **simple** test $\mathcal{T}$ deciding on $H_1, \dots, H_L$ can be lower-bounded as

$$\text{Risk}[\mathcal{T}|H_1, \dots, H_L] \geq \frac{1}{2} \int_{\Omega} \min [P_\ell(d\omega), P_{\ell'}(d\omega)],$$

where, by convention, the integral in the right hand side is

$$\int_{\Omega} \min[p_\ell(\omega), p_{\ell'}(\omega)] \Pi(d\omega),$$

with $p_\ell, p_{\ell'}$ being the densities of $P_\ell, P_{\ell'}$ w.r.t. $\Pi = P_\ell + P_{\ell'}$.

Indeed, risk of $\mathcal{T}$ cannot be less than the risk of the naturally induced by $\mathcal{T}$ simple test deciding on two simple hypotheses $P = P_\ell, P = P_{\ell'}$, specifically, the simple test which, given observation ω accepts the hypothesis $P = P_\ell$ whenever $\ell \in \mathcal{T}(\omega)$ and accepts the hypothesis $P = P_{\ell'}$ otherwise.

Extension: Euclidean separation

Given observation $\omega = x + \xi$ with observation noise $\xi \sim \mathcal{N}(0, I)$, we want to decide on two *composite* hypotheses H_1 , and H_2 :

$$H_1 : x \in X_1, \quad H_2 : x \in X_2,$$

where X_1, X_2 are nonempty *nonintersecting, closed and convex* sets, and one of the sets is bounded.

Elementary fact: *With X_1, X_2 as above, consider the convex minimization problem*

$$\text{Opt} = \min_{x^1 \in X_1, x^2 \in X_2} \frac{1}{2} \|x^1 - x^2\|_2.$$

The problem is solvable. Let (x_^1, x_*^2) be an optimal solution, and let*

$$\phi(\omega) = e^T \omega - c, \quad e = \frac{x_*^1 - x_*^2}{\|x_*^1 - x_*^2\|_2}, \quad c = \frac{1}{2} e^T [x_*^1 + x_*^2]$$

Then the stripe $\{\omega : -\text{Opt} \leq \phi(\omega) \leq \text{Opt}\}$ separates X_1 and X_2 :

$$\phi(x^1) \geq \phi(x_*^1) = \text{Opt} \quad \forall x^1 \in X_1, \quad \phi(x^2) \leq \phi(x_*^2) = -\text{Opt} \quad \forall x^2 \in X_2$$

$$\phi(\omega) = e^T \omega - \frac{1}{2} e^T [x_*^1 + x_*^2], \quad e = \frac{1}{2\text{Opt}} [x_*^1 - x_*^2]$$

where $[x_*^1, x_*^2]$ is an optimal solution to convex optimization problem

$$\text{Opt} = \min_{x^1 \in X_1, x^2 \in X_2} \frac{1}{2} \|x^1 - x^2\|_2$$

While optimal solution is not necessarily unique, $\phi(\cdot)$ is *uniquely defined* by X_1, X_2 .

- Given $\delta_1 \geq 0, \delta_2 \geq 0$ with $\delta_1 + \delta_2 = 2\text{Opt}$, $\phi(\cdot)$ specifies simple *Euclidean Separation Test* $\mathcal{T}$ induced by $X_1, X_2, \delta_1, \delta_2$:

$$\mathcal{T}(\omega) = \begin{cases} \{1\}, & \phi(\omega) \geq \frac{1}{2}[\delta_2 - \delta_1] \\ \{2\}, & \text{otherwise} \end{cases}$$

Observation: Let $\xi \sim \mathcal{N}(0, I)$. Given observation $\omega = x + \xi$ the Euclidean Separation Test $\mathcal{T}$ decides on the hypotheses

$$H_1 : x \in X_1, \quad H_2 : x \in X_2$$

with risks satisfying

$$\begin{aligned} \text{Risk}_1[\mathcal{T}|H_1, H_2] &\leq \text{Prob}_{\zeta \sim N(0,1)} \{\zeta \geq \delta_1\} = \frac{1}{2}\text{erfc}(\delta_1/\sqrt{2}), \\ \text{Risk}_2[\mathcal{T}|H_1, H_2] &\leq \text{Prob}_{\zeta \sim N(0,1)} \{\zeta \geq \delta_2\} = \frac{1}{2}\text{erfc}(\delta_2/\sqrt{2}). \end{aligned}$$

In addition, when $\delta_1 = \delta_2 = \text{Opt}$, $\mathcal{T}$ is the minimum risk test deciding on H_1 , H_2 . The risk of this test is

$$\text{Risk}[\mathcal{T}|H_1, H_2] = \text{Prob}_{\zeta \sim N(0,1)} \{\zeta \geq \text{Opt}\} = \frac{1}{2}\text{erfc}(\text{Opt}/\sqrt{2}).$$

Note: By Neyman-Pearson lemma, this risk is the smallest attainable under circumstances.

From Pairwise to Multiple Hypotheses Testing

Situation: We are given L families of probability distributions $\mathcal{P}_\ell$, $1 \leq \ell \leq L$, on observation space Ω , and observe a realization of random variable $\omega \sim P$ taking values in Ω . Given ω , we want to decide on L hypotheses

$$H_\ell : P \in \mathcal{P}_\ell, \quad 1 \leq \ell \leq L.$$

Our ideal goal would be to find a low-risk simple test deciding on these hypotheses.

Note: It may happen that the “ideal goal” is not achievable, for instance, when some pairs of families $\mathcal{P}_\ell$ have nonempty intersections. When $\mathcal{P}_\ell \cap \mathcal{P}_{\ell'} \neq \emptyset$ for some $\ell \neq \ell'$, there is no way to decide on the hypotheses with risk $< 1/2$.

Fortunately, impossibility to decide reliably on all L hypotheses “individually” does not mean that no meaningful inferences can be done.

*When handling multiple hypotheses which cannot be reliably decided upon “as they are,” it makes sense to speak about **testing the hypotheses “up to closeness.”***

$$\omega \sim P, \quad H_\ell : P \in \mathcal{P}_\ell, \quad 1 \leq \ell \leq L$$

Closeness relation $\mathcal{C}$ on L hypotheses $H_1, \dots, H_L$ is defined as a set of pairs (ℓ, ℓ') with $1 \leq \ell, \ell' \leq L$; we interpret the relation $(\ell, \ell') \in \mathcal{C}$ as the fact that the hypotheses H_ℓ and $H_{\ell'}$ are close to each other.

We always assume that

- $\mathcal{C}$ contains all “diagonal pairs” (ℓ, ℓ) , $1 \leq \ell \leq L$: (“every hypothesis is close to itself”)
- $(\ell, \ell') \in \mathcal{C}$ if and only if $(\ell', \ell) \in \mathcal{C}$ (“closeness is symmetric”)

Note: By symmetry of $\mathcal{C}$, the relation $(\ell, \ell') \in \mathcal{T}$ is in fact a property of *un-ordered* pair $\{\ell, \ell'\}$.

“Up to closeness” risks

Let $\mathcal{T}$ be a test deciding on $H_1, \dots, H_L$; given observation ω , $\mathcal{T}$ accepts all hypotheses H_ℓ with indexes $\ell \in \mathcal{T}(\omega)$ and rejects all other hypotheses.

We say that ℓ -th partial $\mathcal{C}$ -risk of test $\mathcal{T}$ is $\leq \epsilon$, if *whenever H_ℓ is true: $\omega \sim P \in \mathcal{P}_\ell$, the P -probability of the event*

$\mathcal{T}$ accepts H_ℓ : $\ell \in \mathcal{T}(\omega)$ and all hypotheses $H_{\ell'}$ accepted by $\mathcal{T}$ are $\mathcal{C}$ -close to H_ℓ : $(\ell, \ell') \in \mathcal{C} \forall \ell' \in \mathcal{T}(\omega)$

is at least $1 - \epsilon$.

ℓ -th partial $\mathcal{C}$ -risk of $\mathcal{T}$ is the smallest ϵ with the outlined property:

$$\text{Risk}_\ell^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L] = \sup_{P \in \mathcal{P}_\ell} \text{Prob}_{\omega \sim P} \left\{ [\ell \notin \mathcal{T}(\omega)] \text{ or } [\exists \ell' \in \mathcal{T}(\omega) : (\ell, \ell') \notin \mathcal{C}] \right\}$$

$\mathcal{C}$ -risk of $\mathcal{T}$ is the largest of the partial $\mathcal{C}$ -risks of the test:

$$\text{Risk}^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L] = \max_{1 \leq \ell \leq L} \text{Risk}_\ell^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L].$$

Multiple Hypothesis Testing via Pairwise Tests

Assume that for every *unordered* pair $\{\ell, \ell'\}$ with $(\ell, \ell') \notin \mathcal{C}$ we are given a *simple* test $\mathcal{T}_{\{\ell, \ell'\}}$ deciding on H_ℓ vs. $H_{\ell'}$ via observation ω .

Construction:

- For $(\ell, \ell') \notin \mathcal{C}$, so that $\ell \neq \ell'$, we define function $T_{\ell\ell'}(\omega)$ as follows:

$$T_{\ell\ell'}(\omega) = \begin{cases} 1, & \mathcal{T}_{\{\ell, \ell'\}}(\omega) = \{\ell\} \\ -1, & \mathcal{T}_{\{\ell, \ell'\}}(\omega) = \{\ell'\} \end{cases}.$$

Since $\mathcal{T}_{\{\ell, \ell'\}}$ is a simple test, $T_{\ell\ell'}(\cdot)$ is well defined and takes values ± 1 .

- For $(\ell, \ell') \in \mathcal{C}$, we set $T_{\ell\ell'}(\cdot) \equiv 0$.

By construction, we have $T_{\ell\ell'}(\omega) \equiv -T_{\ell'\ell}(\omega)$, $1 \leq \ell, \ell' \leq L$.

- The test $\mathcal{T}$ is as follows: *given observation ω , we build $L \times L$ matrix $T(\omega) = [T_{\ell\ell'}(\omega)]$ and accept exactly those of the hypotheses H_ℓ for which ℓ -th row in $T(\omega)$ is nonnegative*, that is, all tests $\mathcal{T}_{\{\ell, \ell'\}}$ with $(\ell, \ell') \notin \mathcal{C}$ accept H_ℓ , observation being ω .

Example: $L = 4$, $\mathcal{C} = \{(1, 1), (2, 2), (3, 3), (4, 4), \{1, 2\}, \{2, 3\}, \{3, 4\}\}$

Suppose that given tests $\mathcal{T}_{\{1,3\}}$, $\mathcal{T}_{\{1,4\}}$, $\mathcal{T}_{\{2,4\}}$ and observation ω :

- $\mathcal{T}_{\{1,3\}}$ accepts H_1 , $\mathcal{T}_{\{1,4\}}$ accepts H_1 , $\mathcal{T}_{\{2,4\}}$ accepts H_4 , so we get

$$T(\omega) = \begin{bmatrix} 0 & 0 & +1 & +1 \\ 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 \\ -1 & +1 & 0 & 0 \end{bmatrix}$$

$\Rightarrow$ Aggregated test $\mathcal{T}$ accepts H_1

- when $\mathcal{T}_{\{1,3\}}$ accepts H_1 , $\mathcal{T}_{\{1,4\}}$ accepts H_1 , $\mathcal{T}_{\{2,4\}}$ accepts H_2 , we get

$$T(\omega) = \begin{bmatrix} 0 & 0 & +1 & +1 \\ 0 & 0 & 0 & +1 \\ -1 & 0 & 0 & 0 \\ -1 & -1 & 0 & 0 \end{bmatrix}$$

$\Rightarrow$ Aggregated test $\mathcal{T}$ accepts H_1 and H_2

Observation: When $\mathcal{T}$ accepts some hypothesis H_ℓ , *all* hypotheses accepted by $\mathcal{T}$ are $\mathcal{C}$ -close to H_ℓ .

Indeed, if ℓ -th row in $T(\omega)$ is nonnegative and ℓ' is *not* $\mathcal{C}$ -close to ℓ , we have

$$\begin{aligned} & T_{\ell\ell'}(\omega) \geq 0 \text{ and } T_{\ell\ell'}(\omega) \in \{-1, 1\} \\ \Rightarrow & T_{\ell\ell'}(\omega) = 1 \\ \Rightarrow & T_{\ell'\ell}(\omega) = -T_{\ell\ell'}(\omega) = -1 \\ \Rightarrow & \ell'\text{-th row in } T(\omega) \text{ is *not* nonnegative} \Rightarrow \ell' \text{ is *not* accepted} \end{aligned}$$

Risk analysis. For $(\ell, \ell') \notin \mathcal{C}$, let

$$\begin{aligned} \epsilon_{\ell\ell'} &= \text{Risk}_1[\mathcal{T}_{\{\ell, \ell'\}} | H_\ell, H_{\ell'}] = \sup_{P \in \mathcal{P}_\ell} \text{Prob}_{\omega \sim P} \{\ell \notin \mathcal{T}_{\{\ell, \ell'\}}(\omega)\} \\ &= \sup_{P \in \mathcal{P}_\ell} \text{Prob}_{\omega \sim P} \{T_{\ell\ell'}(\omega) = -1\} = \sup_{P \in \mathcal{P}_\ell} \text{Prob}_{\omega \sim P} \{T_{\ell'\ell}(\omega) = 1\} \\ &= \sup_{P \in \mathcal{P}_\ell} \text{Prob}_{\omega \sim P} \{\ell' \in \mathcal{T}_{\{\ell, \ell'\}}(\omega)\} = \text{Risk}_2[\mathcal{T}_{\{\ell, \ell'\}} | H_{\ell'}, H_\ell]. \end{aligned}$$

Proposition. *One has*

$$\text{Risk}_\ell^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L] \leq \epsilon_\ell := \sum_{\ell': (\ell, \ell') \notin \mathcal{C}} \epsilon_{\ell\ell'}.$$

Indeed, let us fix ℓ , and **let H_ℓ be true**. Let $P \in \mathcal{P}_\ell$ be the distribution of observation ω , and let $I = \{\ell' \leq L : (\ell, \ell') \notin \mathcal{C}\}$. For $\ell' \in I$, let $\mathcal{E}_{\ell'}$ be the event $\{\omega : T_{\ell\ell'}(\omega) = -1\}$. We have $\text{Prob}_{\omega \sim P}(\mathcal{E}_{\ell'}) \leq \epsilon_{\ell\ell'}$ (by definition of $\epsilon_{\ell\ell'}$) $\Rightarrow \text{Prob}_{\omega \sim P}(\underbrace{\bigcup_{\ell' \in I} \mathcal{E}_{\ell'}}_{\mathcal{E}}) \leq \epsilon_\ell$.

When the event $\mathcal{E}$ does *not* take place, we have $T_{\ell\ell'}(\omega) = 1$ for all $\ell' \in I$

$\Rightarrow T_{\ell\ell'}(\omega) \geq 0$ for all $\ell', 1 \leq \ell' \leq L$

$\Rightarrow \ell \in \mathcal{T}(\omega)$

$\Rightarrow$ (by Observation) $\{\ell \in \mathcal{T}(\omega)\} \& \{(\ell, \ell') \in \mathcal{C} \forall \ell' \in \mathcal{T}(\omega)\}$.

By definition of partial $\mathcal{C}$ -risk, we get

$$\text{Risk}_\ell^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L] \leq \text{Prob}_{\omega \sim P}(\mathcal{E}) \leq \epsilon_\ell.$$

□

Testing Multiple Hypotheses via Euclidean Separation

Situation: We are given L nonempty, closed and bounded convex sets $X_\ell \subset \mathbb{R}^d$, $1 \leq \ell \leq L$ and an observation $\omega = x + \xi$, so that

- x is a deterministic sequence of signals, and $\xi \sim \mathcal{N}(0, I)$.

Given ω^K , we want to decide up to closeness $\mathcal{C}$ on L hypotheses $H_\ell : x \in X_\ell$.

Equivalently: Sets $X_\ell \subset \mathbb{R}^d$ specify L families of distributions

$$\mathcal{P}_\ell = \{\mathcal{N}(x, I) : x \in X_\ell\}, \quad 1 \leq \ell \leq L$$

Given ω , we want to decide, up to closeness $\mathcal{C}$, on L hypotheses

$$H_\ell : P \in \mathcal{P}_\ell, \quad 1 \leq \ell \leq L$$

on the distribution P^K of observation ω^K .

We intend to assemble pairwise Euclidean separation tests.

Standing Assumption: *Whenever ℓ, ℓ' are not $\mathcal{C}$ -close: $(\ell, \ell') \notin \mathcal{C}$, the sets $X_\ell, X_{\ell'}$ do not intersect.*

Building blocks. For $(\ell, \ell') \notin \mathcal{C}$, we solve convex optimization problems

$$\text{Opt}_{\ell\ell'} = \min_{u \in X_\ell, v \in X_{\ell'}} \frac{1}{2} \|u - v\|_2. \quad (P_{\ell\ell'})$$

By Standing Assumption, $\text{Opt}_{\ell\ell'} > 0$. Optimal solution (u_*, v_*) to $(P_{\ell\ell'})$ defines affine functions

$$\phi_{\ell\ell'}(\omega) = e_{\ell\ell'}^T \omega - c_{\ell\ell'} \quad e_{\ell\ell'} = \frac{u_* - v_*}{\|u_* - v_*\|_2}, \quad c_{\ell\ell'} = \frac{1}{2} e_{\ell\ell'}^T [u_* + v_*]$$

Note: We have $\phi_{\ell\ell'}(\cdot) \equiv -\phi_{\ell'\ell}(\cdot)$ for all $(\ell, \ell') \notin \mathcal{C}$.

As we know, whenever $\delta_{\ell\ell'} \geq 0, \delta_{\ell'\ell} \geq 0$ satisfy

$$2\text{Opt}_{\ell\ell'} = \delta_{\ell\ell'} + \delta_{\ell'\ell}$$

it holds

$$\begin{aligned} \forall u \in X_\ell : \quad \text{Prob}_u \{ \phi(u + \xi) < \tfrac{1}{2} [\delta_{\ell'\ell} - \delta_{\ell\ell'}] \} &\leq \Gamma(\delta_{\ell\ell'}) := \int_{\delta_{\ell\ell'}}^{\infty} \gamma(s) ds, \\ \forall v \in X_{\ell'} : \quad \text{Prob}_v \{ \phi(v + \xi) \geq \tfrac{1}{2} [\delta_{\ell'\ell} - \delta_{\ell\ell'}] \} &\leq \Gamma(\delta_{\ell'\ell}) := \int_{\delta_{\ell'\ell}}^{\infty} \gamma(s) ds \end{aligned}$$

where $\gamma(\cdot)$ is the standard normal density.

$$\begin{aligned}
\ell, \ell' : (\ell, \ell') \notin \mathcal{C} &\Rightarrow \text{Opt}_{\ell\ell'} = \min_{u \in X_\ell, v \in X_{\ell'}} \frac{1}{2} \|u - v\|_2 > 0 = \text{Opt}_{\ell'\ell} \\
\Rightarrow u_*, v_*, \quad \phi_{\ell\ell'}(\omega) = e_{\ell\ell'}^T \omega - c_{\ell\ell'} &\equiv -\phi_{\ell'\ell}(\omega) \quad \left[e_{\ell\ell'} = \frac{u_* - v_*}{\|u_* - v_*\|_2}, \quad c_{\ell\ell'} = \frac{1}{2} e_{\ell\ell'}^T [u_* + v_*] \right] \\
\delta_{\ell\ell'} \geq 0, \quad \delta_{\ell'\ell} \geq 0, \quad 2\text{Opt}_{\ell\ell'} &= \delta_{\ell\ell'} + \delta_{\ell'\ell} \quad (*) \\
\forall u \in X_\ell : \quad \text{Prob}_{\xi \sim P} \left\{ \phi(u + \xi) < \frac{1}{2} [\delta_{\ell'\ell} - \delta_{\ell\ell'}] \right\} &\leq \Gamma(\delta_{\ell\ell'}) := \int_{\delta_{\ell\ell'}}^{\infty} \gamma(s) ds \\
\forall v \in X_{\ell'} : \quad \text{Prob}_{\xi \sim P} \left\{ \phi(v + \xi) \geq \frac{1}{2} [\delta_{\ell'\ell} - \delta_{\ell\ell'}] \right\} &\leq \Gamma(\delta_{\ell'\ell}) := \int_{\delta_{\ell'\ell}}^{\infty} \gamma(s) ds \quad (!)
\end{aligned}$$

Assembling building blocks

- For ℓ, ℓ' with $(\ell, \ell') \notin \mathcal{C}$ we select $\delta_{\ell\ell'}$ satisfying (*), thus arriving at simple tests

$$\mathcal{T}_{\{\ell, \ell'\}}(\omega) = \begin{cases} \{\ell\}, & \phi_{\ell\ell'}(\omega) \geq \frac{1}{2} [\delta_{\ell'\ell} - \delta_{\ell\ell'}] \\ \{\ell'\}, & \phi_{\ell\ell'}(\omega) < \frac{1}{2} [\delta_{\ell'\ell} - \delta_{\ell\ell'}] \end{cases}$$

- Further, we use our general construction to assemble tests $\{\mathcal{T}_{\{\ell, \ell'\}} : (\ell, \ell') \notin \mathcal{C}\}$ into test $\mathcal{T}$ deciding on $H_1, \dots, H_L$

Note: By (!), associated with tests $\mathcal{T}_{\{\ell, \ell'\}}$ quantities $\epsilon_{\ell\ell'}$ satisfy the relations

$$\epsilon_{\ell\ell'} \leq \Gamma(\delta_{\ell\ell'}) := \int_{\delta_{\ell\ell'}}^{\infty} \gamma(s) ds, \text{ whence } \text{Risk}_{\ell}^{\mathcal{C}}[\mathcal{T} | H_1, \dots, H_L] \leq \sum_{\ell' : (\ell, \ell') \notin \mathcal{C}} \Gamma(\delta_{\ell\ell'}).$$

$$\begin{aligned}
\ell, \ell' : (\ell, \ell') \notin \mathcal{C} &\Rightarrow \text{Opt}_{\ell\ell'} = \min_{u \in X_\ell, v \in X_{\ell'}} \frac{1}{2} \|u - v\|_2 \\
&\Rightarrow \delta_{\ell\ell'} \geq 0, \delta_{\ell'\ell} \geq 0, 2\text{Opt}_{\ell\ell'} = \delta_{\ell\ell'} + \delta_{\ell'\ell} \\
\Rightarrow \mathcal{T} : \text{Risk}_\ell^{\mathcal{C}}[\mathcal{T} | H_1, \dots, H_L] &\leq \sum_{\ell' : (\ell, \ell') \notin \mathcal{C}} \Gamma(\delta_{\ell\ell'}), \Gamma(\delta) = \int_{\delta}^{\infty} \gamma(s) ds
\end{aligned}$$

Refining the construction.

We can try to optimize the “profile” of (upper bounds on) partial $\mathcal{C}$ -risks of the test $\mathcal{T}$ over the “free parameters” $\delta_{\ell\ell'}$, $(\ell, \ell') \notin \mathcal{C}$, of the construction.

A model here may be as follows: given nonnegative *weight matrix* W and nonnegative vectors α and β , we want to minimize “scale factor” t under the constraint

$$W \text{vec}[\text{Risk}_\ell^{\mathcal{C}}[\mathcal{T} | H_1, \dots, H_L]]_{\ell=1}^L \leq \alpha + t\beta$$

This problem can be safely approximated by the optimization problem

$$\min_{\{\delta_{\ell\ell'}\}, t} \left\{ t : \begin{aligned} &W \text{vec} \left[\sum_{\ell' : (\ell, \ell') \notin \mathcal{C}} \Gamma(\delta_{\ell\ell'}) \right]_{\ell=1}^L \leq \alpha + t\beta \\ &\delta_{\ell\ell'} \geq 0, \delta_{\ell\ell'} + \delta_{\ell'\ell} = 2\text{Opt}_{\ell\ell'}, (\ell, \ell') \notin \mathcal{C} \end{aligned} \right\} \quad (\#)$$

Note: When $\gamma(\cdot)$ is nonincreasing on $\mathbb{R}_+$ (as is the case for Gaussian density), function $\Gamma(\delta) = \int_{\delta}^{\infty} \gamma(s) ds$ is convex on $\mathbb{R}_+$

$\Rightarrow (\#)$ is an explicit Convex Programming problem!

HYPOTHESIS TESTING, II

- *Detector-Based Tests*
- *Detectors & Detector-Based Pairwise Tests*
- *Testing “up to Closeness”*
- *Simple Observation Schemes*
 - *Minimum Risk Detectors*
 - *Near-Optimal Tests*
- *Applications*
 - *Recovery of Functionals*
 - *Sequential Hypothesis Testing*
 - *Measurement Design*

Detectors & Detector-Based Pairwise Tests

Situation: Given two families $\mathcal{P}_1, \mathcal{P}_2$ of probability distributions on a given observation space Ω and an observation $\omega \sim P$ with P known to belong to $\mathcal{P}_1 \cup \mathcal{P}_2$, we want to decide whether

$P \in \mathcal{P}_1$ (hypothesis H_1) or $P \in \mathcal{P}_2$ (hypothesis H_2).

Detectors. A *detector* is a function $\phi : \Omega \rightarrow \mathbb{R}$. *Risks* of a detector ϕ w.r.t. $\mathcal{P}_1$ and $\mathcal{P}_2$ are defined as

$$\begin{aligned}\text{Risk}_1[\phi|\mathcal{P}_1, \mathcal{P}_2] &= \sup_{P \in \mathcal{P}_1} \int_{\Omega} e^{-\phi(\omega)} P(d\omega), \\ \text{Risk}_2[\phi|\mathcal{P}_1, \mathcal{P}_2] &= \sup_{P \in \mathcal{P}_2} \int_{\Omega} e^{\phi(\omega)} P(d\omega) \\ \text{Risk}_1[\phi|\mathcal{P}_1, \mathcal{P}_2] &= \text{Risk}_2[-\phi|\mathcal{P}_2, \mathcal{P}_1]\end{aligned}$$

Simple test $\mathcal{T}_\phi$ associated with detector ϕ , given observation ω ,

- accepts H_1 when $\phi(\omega) \geq 0$
- accepts H_2 when $\phi(\omega) < 0$.

Immediate observation:

$$\text{Risk}_1[\mathcal{T}_\phi|H_1, H_2] \leq \text{Risk}_1[\phi|\mathcal{P}_1, \mathcal{P}_2]$$

$$\text{Risk}_2[\mathcal{T}_\phi|H_1, H_2] \leq \text{Risk}_2[\phi|\mathcal{P}_1, \mathcal{P}_2]$$

Indeed, $\text{Prob}_{\omega \sim P} \{\omega : \psi(\omega) \geq 0\} \leq \int e^{\psi(\omega)} P(d\omega)$.

$$\begin{aligned}\text{Risk}_1[\phi|\mathcal{P}_1, \mathcal{P}_2] &= \sup_{P \in \mathcal{P}_1} \int_{\Omega} e^{-\phi(\omega)} P(d\omega), \\ \text{Risk}_2[\phi|\mathcal{P}_1, \mathcal{P}_2] &= \sup_{P \in \mathcal{P}_2} \int_{\Omega} e^{\phi(\omega)} P(d\omega)\end{aligned}$$

Elementary Calculus of Detectors

- Renormalization (shift): $\phi(\cdot) \Rightarrow \phi_a(\cdot) = \phi(\cdot) - a$

$$\Rightarrow \begin{cases} \text{Risk}_1[\phi_a|\mathcal{P}_1, \mathcal{P}_2] = e^a \text{Risk}_1[\phi|\mathcal{P}_1, \mathcal{P}_2] \\ \text{Risk}_2[\phi_a|\mathcal{P}_1, \mathcal{P}_2] = e^{-a} \text{Risk}_2[\phi|\mathcal{P}_1, \mathcal{P}_2] \end{cases}$$

$\Rightarrow$ *What matters, is the product*

$$[\text{Risk}[\phi|\mathcal{P}_1, \mathcal{P}_2]]^2 := \text{Risk}_1[\phi|\mathcal{P}_1, \mathcal{P}_2] \text{Risk}_2[\phi|\mathcal{P}_1, \mathcal{P}_2]$$

*of partial risks of a detector. Shifting detector by a constant, we can redistribute this product between factors as we want, e.g., make the detector **balanced**:*

$$\text{Risk}[\phi|\mathcal{P}_1, \mathcal{P}_2] = \text{Risk}_1[\phi|\mathcal{P}_1, \mathcal{P}_2] = \text{Risk}_2[\phi|\mathcal{P}_1, \mathcal{P}_2].$$

Detectors are well-suited to passing to multiple observations.

For $1 \leq k \leq K$, let

- $\mathcal{P}_{1,k}, \mathcal{P}_{2,k}$ be families of probability distributions on observation spaces Ω_k ,
- ϕ_k be detectors on Ω_k .

Families $\{\mathcal{P}_{1,k}, \mathcal{P}_{2,k}\}_{k=1}^K$ give rise to families of product distributions

$$\mathcal{P}_{\chi}^{\oplus, K} = \bigoplus_{k=1}^K \mathcal{P}_{\chi, k} := \{P^K = P_1 \times \dots \times P_K : P_k \in \mathcal{P}_{\chi, k}, 1 \leq k \leq K\}.$$

on $\Omega^K = \Omega_1 \times \dots \times \Omega_K$, $\chi = 1, 2$. We call $\mathcal{P}_{\chi}^K$ *direct products* of $\mathcal{P}_{\chi, k}$.

Detectors $\phi_1, \dots, \phi_K$ give rise to detector ϕ^K on Ω^K :

$$\phi^K(\underbrace{\omega_1, \dots, \omega_K}_{\omega^K}) = \sum_{k=1}^K \phi_k(\omega_k).$$

Observation: For $\chi = 1, 2$, we have

$$\text{Risk}_{\chi}[\phi^K | \mathcal{P}_1^{\oplus, K}, \mathcal{P}_2^{\oplus, K}] = \prod_{k=1}^K \text{Risk}_{\chi}[\phi_k | \mathcal{P}_{1,k}, \mathcal{P}_{2,k}].$$

From pairwise detectors to detectors for unions. Assume that we are given an observation space Ω along with

- R families $\mathcal{R}_r, r = 1, \dots, R$ of “red” probability distributions on Ω ,
- B families $\mathcal{B}_b, b = 1, \dots, B$ of “brown” probability distributions on Ω ,
- pairwise detectors $\phi_{rb}(\cdot), 1 \leq r \leq R, 1 \leq b \leq B$.

Let us aggregate the red and the brown families into 2 groups as follows

$$\mathcal{R} = \bigcup_{r=1}^R \mathcal{R}_r, \quad \mathcal{B} = \bigcup_{b=1}^B \mathcal{B}_b$$

and assemble detectors ϕ_{rb} into a single detector

$$\phi(\omega) = \max_{r \leq R} \min_{b \leq B} \phi_{rb}(\omega).$$

- *We have*

$$\text{Risk}_1[\phi|\mathcal{R}, \mathcal{B}] \leq \max_{r \leq R} \sum_{b \leq B} \text{Risk}_1[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b],$$

$$\text{Risk}_2[\phi|\mathcal{R}, \mathcal{B}] \leq \max_{b \leq B} \sum_{r \leq R} \text{Risk}_2[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b].$$

$$\begin{aligned}\text{Risk}_1[\phi|\mathcal{R}, \mathcal{B}] &\leq \max_{r \leq R} \sum_{b \leq B} \text{Risk}_1[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b], \\ \text{Risk}_2[\phi|\mathcal{R}, \mathcal{B}] &\leq \max_{b \leq B} \sum_{r \leq R} \text{Risk}_2[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b],\end{aligned}$$

Indeed, because $-\max_r \min_b \phi_{rb}(\omega) = \min_r \max_b -\phi_{rb}(\omega)$,

$$\begin{aligned}P \in \mathcal{R}_{r_*} \Rightarrow \int e^{-\phi(\omega)} P(d\omega) &= \int e^{\min_r \max_b [-\phi_{rb}(\omega)]} P(d\omega) \\ &\leq \int e^{\max_b [-\phi_{r_*b}(\omega)]} P(d\omega) \leq \sum_b \int e^{-\phi_{r_*b}(\omega)} P(d\omega) \\ &\leq \sum_b \text{Risk}_1[\phi_{r_*b}|\mathcal{R}_{r_*}, \mathcal{B}_b] \\ \Rightarrow \text{Risk}_1[\phi|\mathcal{R}, \mathcal{B}] &\leq \max_{r \leq R} \sum_{b \leq B} \text{Risk}_1[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b].\end{aligned}$$

In the same way,

$$\begin{aligned}P \in \mathcal{B}_{b_*} \Rightarrow \int e^{\phi(\omega)} P(d\omega) &= \int e^{\max_r \min_b [\phi_{rb}(\omega)]} P(d\omega) \leq \int e^{\max_r \phi_{rb_*}(\omega)} P(d\omega) \\ &\leq \sum_r \int e^{\phi_{rb_*}(\omega)} P(d\omega) \leq \sum_r \text{Risk}_2[\phi_{rb_*}|\mathcal{R}_r, \mathcal{B}_{b_*}] \\ \Rightarrow \text{Risk}_2[\phi|\mathcal{R}, \mathcal{B}] &\leq \max_{b \leq B} \sum_{r \leq R} \text{Risk}_2[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b].\end{aligned}$$

Refinement: W.l.o.g. we may assume that the detectors ϕ_{rb} are balanced:

$$\epsilon_{rb} := \text{Risk}[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b] = \text{Risk}_1[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b] = \text{Risk}_2[\phi_{rb}|\mathcal{R}_r, \mathcal{B}_b].$$

Consider matrices

$$E = \begin{bmatrix} \epsilon_{1,1} & \cdots & \epsilon_{1,B} \\ \vdots & \cdots & \vdots \\ \epsilon_{R,1} & \cdots & \epsilon_{R,B} \end{bmatrix}, \quad F = \left[\begin{array}{c|c} & E \\ \hline E^T & \end{array} \right]$$

Note: The maximal eigenvalue θ of F is the spectral norm $\|E\|_{2,2}$ of E , and the leading eigenvector $[g; f]$ can be selected to be positive (*Perron-Frobenius theorem*), and we have

$$\theta g = E f \quad \theta f = E^T g$$

- Let us pass from the detectors ϕ_{rb} to shifted detectors $\psi_{rb} = \phi_{rb} - \ln(f_b/g_r)$ and assemble the shifted detectors into the detector

$$\psi(\omega) = \max_{r \leq R} \min_{b \leq B} \psi_{rb}(\omega)$$

By the previous observation

$$\begin{aligned} \text{Risk}_1[\psi|\mathcal{R}, \mathcal{B}] &\leq \max_r \sum_b \text{Risk}_1(\psi_{rb}|\mathcal{R}_r, \mathcal{B}_b) = \max_r \sum_b \epsilon_{rb} (f_b/g_r) = \max_r [(E f)_r / g_r] = \theta \\ \text{Risk}_2[\psi|\mathcal{R}, \mathcal{B}] &\leq \max_b \sum_r \text{Risk}_2(\psi_{rb}|\mathcal{R}_r, \mathcal{B}_b) = \max_b \sum_r \epsilon_{rb} (g_r/f_b) = \max_b [(E^T g)_b / f_b] = \theta \\ \Rightarrow \text{Partial risks of detector } \psi \text{ on aggregated families } \mathcal{R}, \mathcal{B} \text{ are } &\leq \theta = \|E\|_{2,2}. \end{aligned}$$

Detector-Based Tests “Up to Closeness”

Situation: We are given

- L families of probability distributions $\mathcal{P}_\ell$, $\ell = 1, \dots, L$, on observation space Ω , giving rise to L hypotheses H_ℓ , on the distribution P of random observation $\omega \in \Omega$:

$$H_\ell : P \in \mathcal{P}_\ell, \quad 1 \leq \ell \leq L;$$

- closeness relation $\mathcal{C}$;
- system of balanced detectors

$$\{\phi_{\ell\ell'} : \ell < \ell', (\ell, \ell') \notin \mathcal{C}\}$$

along with upper bounds $\epsilon_{\ell\ell'}$ on the detector risks:

$$\forall (\ell, \ell' : \ell < \ell', (\ell, \ell') \notin \mathcal{C}) : \begin{cases} \int_{\Omega} e^{-\phi_{\ell\ell'}(\omega)} P(d\omega) \leq \epsilon_{\ell\ell'} & \forall P \in \mathcal{P}_\ell \\ \int_{\Omega} e^{\phi_{\ell\ell'}(\omega)} P(d\omega) \leq \epsilon_{\ell\ell'} & \forall P \in \mathcal{P}_{\ell'} \end{cases}$$

- Our goal is to build single-observation test deciding on hypotheses $H_1, \dots, H_L$ up to closeness $\mathcal{C}$.

Construction:

- We set

$$\phi_{\ell\ell'}(\omega) = \begin{cases} -\phi_{\ell'\ell}(\omega), & \ell > \ell', (\ell, \ell') \notin \mathcal{C}, \\ 0, & (\ell, \ell') \in \mathcal{C}, \end{cases} \quad \epsilon_{\ell\ell'} = \begin{cases} \epsilon_{\ell'\ell}, & \ell > \ell', (\ell, \ell') \notin \mathcal{C}, \\ 1, & (\ell, \ell') \in \mathcal{C}, \end{cases}$$

thus ensuring that

$$\phi_{\ell\ell'}(\cdot) = -\phi_{\ell'\ell}(\cdot), \text{ and } \epsilon_{\ell\ell'} = \epsilon_{\ell'\ell}, \quad 1 \leq \ell, \ell' \leq L$$
$$\int_{\Omega} e^{-\phi_{\ell\ell'}(\omega)} P(d\omega) \leq \epsilon_{\ell\ell'} \quad \forall (P \in \mathcal{P}_{\ell}, 1 \leq \ell, \ell' \leq L)$$

- Given shifts $a_{\ell\ell'} = -a_{\ell'\ell}$, we specify test $\mathcal{T}$ as follows:

Given observation ω , $\mathcal{T}$ accepts all hypotheses H_{ℓ} such that

$$\phi_{\ell\ell'}(\omega) > a_{\ell\ell'} \quad \forall (\ell' : (\ell, \ell') \notin \mathcal{C})$$

and rejects all other hypotheses.

Proposition. *The $\mathcal{C}$ -risk of $\mathcal{T}$ can be upper-bounded as*

$$\text{Risk}^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L] \leq \max_{\ell \leq L} \sum_{\ell': (\ell, \ell') \notin \mathcal{C}} \epsilon_{\ell\ell'} e^{a_{\ell\ell'}}$$

Optimal shifts: Consider the symmetric nonnegative matrix

$$E = [\epsilon_{\ell\ell'} \chi_{\ell\ell'}]_{\ell,\ell'=1}^L, \quad \chi_{\ell\ell'} = \begin{cases} 1, & (\ell, \ell') \notin \mathcal{C} \\ 0, & (\ell, \ell') \in \mathcal{C} \end{cases},$$

and let $\theta = \|E\|_{2,2}$ be the spectral norm of E , or, which is the same under the circumstances, the largest eigenvalue of E .

By the Perron-Frobenius theorem, for every $\theta' > \theta$ there exists a positive vector f such that

$$Ef \leq \theta' f;$$

the same holds true for $\theta' = \theta$, provided the leading eigenvector of E (which always can be selected to be nonnegative) is positive.

Observation: For $\alpha_{\ell\ell'} = \ln(f_{\ell'}/f_{\ell})$, the risk bound from the proposition reads

$$\text{Risk}^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L] \leq \theta'.$$

Thus, assembling the detectors $\phi_{\ell\ell'}$ appropriately, one can get a test with $\mathcal{C}$ -risk arbitrarily close to $\|E\|_{2,2}$.

Utilizing repeated observations. We can apply the above construction to

- K -repeated observation $\omega^K = (\omega_1, \dots, \omega_K)$ in the role of ω ,
- direct products $\mathcal{P}_\ell^{\oplus, K} = \bigoplus_{k=1}^K \mathcal{P}_\ell$ in the role of families $\mathcal{P}_\ell$, and respective hypotheses $H_\ell^{\oplus, K}$ in the role of hypotheses H_ℓ , $\ell = 1, \dots, L$,
- detectors $\phi_{\ell\ell'}^{(K)}(\omega^K) = \sum_{k=1}^K \phi_{\ell\ell'}(\omega_k)$ in the role of detectors $\phi_{\ell\ell'}$, which allows to replace $\epsilon_{\ell\ell'}$ with $\epsilon_{\ell\ell'}^K$.

As a result, we get K -observation test $\mathcal{T}^K$ such that

$$\text{Risk}^{\mathcal{C}}[\mathcal{T}^K | H_1^{\oplus, K}, \dots, H_L^{\oplus, K}] \leq \theta'_K$$

where θ'_K can be made arbitrarily close (under favorable circumstances, even equal) to the quantity

$$\left\| \left[\epsilon_{\ell\ell'}^K \chi_{\ell\ell'} \right]_{\ell, \ell'=1}^L \right\|_{2,2}, \quad \chi_{\ell\ell'} = \begin{cases} 1, & (\ell, \ell') \notin \mathcal{C} \\ 0, & (\ell, \ell') \in \mathcal{C} \end{cases}$$

In particular, in the case when $\epsilon_{\ell\ell'} < 1$ whenever $(\ell, \ell') \notin \mathcal{C}$, we can ensure that the $\mathcal{C}$ -risk of $\mathcal{T}^K$ converges to 0 exponentially fast as $K \rightarrow \infty$.

“Universality” of detector-based tests.

Proposition. *Let $\mathcal{P}_\chi$, $\chi = 1, 2$, be two families of probability distributions on observation space Ω , and let H_χ , $\chi = 1, 2$, be associate hypotheses on the distribution of an observation.*

Assume that there exists a simple deterministic or randomized test $\mathcal{T}$ deciding on H_1, H_2 with risk $\leq \epsilon \in (0, 1/2)$. Then there exists a detector ϕ with

$$\text{Risk}[\phi|\mathcal{P}_1, \mathcal{P}_2] \leq \epsilon_+ := 2\sqrt{\epsilon[1 - \epsilon]} < 1.$$

Proof. 1^o. Indeed, let $\mathcal{T}$ be deterministic, let $\Omega_\chi = \{\omega \in \Omega : \mathcal{T}(\omega) = \{\chi\}\}$, $\chi = 1, 2$, and let

$$\phi(\omega) = \begin{cases} \frac{1}{2} \ln \left(\frac{[1-\epsilon]}{\epsilon} \right), & \omega \in \Omega_1 \\ \frac{1}{2} \ln (\epsilon/[1-\epsilon]), & \omega \in \Omega_2 \end{cases}$$

Then for $P \in \mathcal{P}_1$ and $\epsilon' = \int_{\Omega_2} P(d\omega)$ we have

$$\begin{aligned} \Rightarrow \int e^{-\phi(\omega)} P(d\omega) &= \sqrt{\epsilon/[1-\epsilon]}(1-\epsilon') + \sqrt{[1-\epsilon]/\epsilon} \epsilon' \\ &= \sqrt{\epsilon/[1-\epsilon]} + \underbrace{\left[\sqrt{[1-\epsilon]/\epsilon} - \sqrt{\epsilon/[1-\epsilon]} \right]}_{\geq 0} \underbrace{\epsilon'}_{\leq \epsilon} \\ &\leq \sqrt{\epsilon/[1-\epsilon]} + \left[\sqrt{[1-\epsilon]/\epsilon} - \sqrt{\epsilon/[1-\epsilon]} \right] \epsilon \\ &= 2\sqrt{\epsilon[1-\epsilon]} \end{aligned}$$

In the same way, when $P \in \mathcal{P}_2$ and $\epsilon' = \int_{\Omega_1} P(d\omega)$,

$$\int e^{\phi(\omega)} P(d\omega) = \sqrt{\epsilon/[1-\epsilon]}(1-\epsilon') + \sqrt{[1-\epsilon]/\epsilon} \epsilon' = 2\sqrt{\epsilon[1-\epsilon]}$$

We conclude that

$$\text{Risk}_\chi[\phi|\mathcal{P}_1, \mathcal{P}_2] \leq 2\sqrt{\epsilon[1-\epsilon]}.$$

2°. Now let $\mathcal{T}$ be randomized. Setting $\mathcal{P}_\chi^+ = \{P \times \text{Uniform}[0, 1] : P \in \mathcal{P}_\chi\}$, $\chi = 1, 2$, $\Omega^+ = \Omega \times [0, 1]$.

By the above there exists a bounded detector $\phi_+ : \Omega^+ \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \forall(P \in \mathcal{P}_1) : \int_{\Omega} \left[\int_0^1 e^{-\phi_+(\omega, s)} ds \right] P(d\omega) ds &\leq \bar{\epsilon} = 2\sqrt{\epsilon[1 - \epsilon]}, \\ \forall(P \in \mathcal{P}_2) : \int_{\Omega} \left[\int_0^1 e^{\phi_+(\omega, s)} ds \right] P(d\omega) &\leq \bar{\epsilon}, \end{aligned}$$

whence, setting $\phi(\omega) = \int_0^1 \phi(\omega, s) ds$ and applying Jensen's Inequality,

$$\forall(P \in \mathcal{P}_1) : \int_{\Omega} e^{-\phi(\omega)} P(d\omega) \leq \bar{\epsilon}, \quad \forall(P \in \mathcal{P}_2) : \int_{\Omega} e^{\phi(\omega)} P(d\omega) \leq \bar{\epsilon} \quad \square$$

- Risk $2\sqrt{\epsilon[1 - \epsilon]}$ of the detector-based test induced by simple test $\mathcal{T}$ is “much worse” than the risk ϵ of $\mathcal{T}$.
- *However, when repeated observations are allowed, we can compensate for risk deterioration $\epsilon \rightarrow 2\sqrt{\epsilon[1 - \epsilon]}$ by passing in the detector-based test from a single observation to a moderate number of them.*

Minimal risk detectors

$$\inf_{\phi} \left\{ \text{Risk}[\phi | \mathcal{P}_1, \mathcal{P}_2] = \min \left\{ \epsilon : \begin{array}{l} \int_{\Omega} e^{-\phi(\omega)} P(d\omega) \leq \epsilon \quad \forall (P \in \mathcal{P}_1) \\ \int_{\Omega} e^{\phi(\omega)} P(d\omega) \leq \epsilon \quad \forall (P \in \mathcal{P}_2) \end{array} \right\} \right\} \quad (!)$$

- Optimization problem specifying risk has constraints *convex* in (ϕ, ϵ)
- When passing from families $\mathcal{P}_{\chi}$, $\chi = 1, 2$, to their convex hulls, the risk of a detector remains intact.

It would be nice to be able to solve (!), thus arriving at the lowest risk detector-based tests.

However: (!) is an optimization problem with *infinite-dimensional* decision “vector” and *infinitely many* constraints $\Rightarrow$ *in general, (!) is intractable.*

Simple observation schemes: Special cases where (!) is efficiently solvable via Convex Optimization.

Simple Observation Schemes

Simple Observation Scheme is a collection

$$\mathcal{O} = ((\Omega, \Pi), \{p_\mu : \mu \in \mathcal{M}\}, \mathcal{F}),$$

where

- (Ω, Π) is a (complete separable metric) *observation space* Ω with (σ -finite σ -additive) *reference measure* Π ,

$$\text{supp } \Pi = \Omega;$$

- $\{p_\mu(\cdot) : \mu \in \mathcal{M}\}$ is a parametric family of probability densities, taken w.r.t. Π , on Ω , and
- $\mathcal{M}$ is a relatively open *convex* set in some $\mathbb{R}^n$
- $p_\mu(\omega)$: *positive* and continuous in $\mu \in \mathcal{M}, \omega \in \Omega$
- $\mathcal{F}$ is a *finite-dimensional* space of continuous functions on Ω containing constants and such that

$$\ln(p_\mu(\cdot)/p_\nu(\cdot)) \in \mathcal{F} \quad \forall \mu, \nu \in \mathcal{M}.$$

- For $\phi \in \mathcal{F}$, the function

$$\mu \mapsto \ln \left(\int_{\Omega} e^{\phi(\omega)} p_\mu(\omega) \Pi(d\omega) \right)$$

is finite and concave in $\mu \in \mathcal{M}$.

Example 1: Gaussian o.s., where (Ω, Π) is $\mathbb{R}^d$ with Lebesgue measure,

$$\{p_\mu(\cdot) = \mathcal{N}(\mu, I_d) : \mu \in \mathbb{R}^d\}, \quad \mathcal{F} = \{\text{affine functions on } \Omega\}$$

$$\Rightarrow \begin{cases} \ln(p_\mu(\cdot)/p_\nu(\cdot)) \in \mathcal{F}, \\ \ln \left(\int_{\Omega} e^{a^T \omega + b} p_\mu(\omega) \Pi(d\omega) \right) = a^T \mu + b + \frac{a^T a}{2} \text{ is concave in } \mu. \end{cases}$$

- Gaussian o.s. is the standard observation model in Signal Processing.

Example 2: Discrete o.s., where (Ω, Π) is finite set $\{1, \dots, d\}$ with counting measure,

$$\left\{ p_\mu(\omega) = \mu_\omega, \mu \in \mathcal{M} = \{\mu > 0 : \sum_{\omega=1}^d \mu_\omega = 1\} \right\}$$

is the set of non-vanishing probability distributions on Ω , and $\mathcal{F} = \{\text{all functions on } \Omega\}$

$$\Rightarrow \begin{cases} \ln(p_\mu(\cdot)/p_\nu(\cdot)) \in \mathcal{F}, \\ \ln \left(\int_{\Omega} e^{\phi(\omega)} p_\mu(\omega) \Pi(d\omega) \right) = \ln \left(\sum_{\omega \in \Omega} e^{\phi(\omega)} \mu_\omega \right) \text{ is concave in } \mu. \end{cases}$$

Example 3: Poisson o.s., where (Ω, Π) , is the nonnegative part $\mathbb{Z}_+^d$ of integer lattice in $\mathbb{R}^d$ equipped with counting measure,

$$\left\{ p_\mu(\omega) = \prod_{i=1}^d \frac{\mu_i^{\omega_i} e^{-\mu_i}}{\omega_i!} : \mu \in \mathcal{M} := \mathbb{R}_{++}^d \right\}$$

is the family of distributions of random vectors with independent across i Poisson entries $\omega_i \sim \text{Poisson}(\mu_i)$, and $\mathcal{F} = \{\text{affine functions on } \Omega\}$

$$\Rightarrow \begin{cases} \ln(p_\mu(\cdot)/p_\nu(\cdot)) \in \mathcal{F}, \\ \ln \left(\int_{\Omega} e^{a^T \omega + b} p_\mu(\omega) \Pi(d\omega) \right) = \sum_i (e^{a_i} - 1) \mu_i \text{ is concave in } \mu. \end{cases}$$

- Poisson o.s. arises in *Poisson Imaging*, including
 - *Positron Emission Tomography*,
 - *Large binocular Telescope*,
 - *Nanoscale Fluorescent Microscopy*.

Example 4: Direct product of simple o.s.'s.

Simple o.s.'s

$$\mathcal{O}_k = \left((\Omega_k, \Pi_k), \{p_{\mu_k, k}(\cdot) : \mu_k \in \mathcal{M}_k\}, \mathcal{F}_k \right), \quad 1 \leq k \leq K$$

give rise to their *direct product* $\bigotimes_{k=1}^K \mathcal{O}_k$ defined as the o.s.

$$\left((\Omega^K, \Pi^K), \{p_{\mu^K}(\cdot) : \mu^K \in \mathcal{M}^K\}, \mathcal{F}^K \right),$$

where

- $\Omega^K = \Omega_1 \times \dots \times \Omega_K, \Pi^K = \Pi_1 \times \dots \times \Pi_K$
- $\mathcal{M}^K = \mathcal{M}_1 \times \dots \times \mathcal{M}_K, p_{(\mu_1, \dots, \mu_K)}(\omega_1, \dots, \omega_K) = \prod_{k=1}^K p_{\mu_k, k}(\omega_k)$
- $\mathcal{F}^K = \{\phi(\underbrace{\omega_1, \dots, \omega_K}_{\omega^K}) = \sum_{k=1}^K \phi_k(\omega_k) : \phi_k \in \mathcal{F}_k, 1 \leq k \leq K\}$
- *A direct product of simple o.s.'s is a simple o.s.*

Example 5: Power of a simple o.s.

When all K o.s.'s in product $\mathcal{O}^K = \bigotimes_{k=1}^K \mathcal{O}_k$ are identical to each other:

$$\mathcal{O}_k = \mathcal{O} := ((\Omega, \Pi), \{p_\mu(\cdot) : \mu \in \mathcal{M}\}, \mathcal{F}), \quad 1 \leq k \leq K,$$

we can “restrict $\mathcal{O}^K$ to its diagonal,” arriving at *K -th power $\mathcal{O}^{(K)}$ of $\mathcal{O}$* :

$$\begin{aligned} \mathcal{O}^{(K)} &= \left((\Omega^K, \Pi^K), \{p_\mu^{(K)}(\cdot) : \mu \in \mathcal{M}\}, \mathcal{F}^{(K)} \right), \\ p_\mu^{(K)}(\omega_1, \dots, \omega_K) &= \prod_{k=1}^K p_\mu(\omega_k), \\ \mathcal{F}^{(K)} &= \{ \phi^{(K)}(\omega^K) = \sum_{k=1}^K \phi(\omega_k) : \phi \in \mathcal{F} \} \end{aligned}$$

- *K -th power of a simple o.s. is a simple o.s.*

$$\varepsilon_*(\mathcal{P}_1, \mathcal{P}_2) = \min_{\phi(\cdot), \epsilon} \left\{ \epsilon : \begin{array}{l} \int_{\Omega} e^{-\phi(\omega)} P(d\omega) \leq \epsilon \forall (P \in \mathcal{P}_1) \\ \int_{\Omega} e^{\phi(\omega)} P(d\omega) \leq \epsilon \forall (P \in \mathcal{P}_2) \end{array} \right\} \quad (!)$$

Main Result. Let $\mathcal{O} = ((\Omega, \Pi), \{p_{\mu}(\cdot) : \mu \in \mathcal{M}\}, \mathcal{F})$ be a simple o.s., and let M_1, M_2 be two nonempty compact convex subsets of $\mathcal{M}$. These subsets give rise to two families of probability distributions $\mathcal{P}_1, \mathcal{P}_2$ on Ω

$$\mathcal{P}_{\chi} = \{P : \text{the density of } P \text{ is } p_{\mu} \text{ with } \mu \in \mathcal{M}_{\chi}\}, \quad \chi = 1, 2,$$

and two hypotheses on the distribution P of random observation $\omega \in \Omega$:

$$H_1 : P \in \mathcal{P}_1 \quad \text{and} \quad H_2 : P \in \mathcal{P}_2$$

Consider the function

$$\Phi(\phi; \mu, \nu) = \frac{1}{2} \left[\ln \left(\int_{\Omega} e^{-\phi(\omega)} p_{\mu}(\omega) \Pi(d\omega) \right) + \ln \left(\int_{\Omega} e^{\phi(\omega)} p_{\nu}(\omega) \Pi(d\omega) \right) \right] : \\ \mathcal{F} \times [M_1 \times M_2] \rightarrow \mathbb{R}.$$

Then

- $\Phi(\phi; \mu, \nu)$ is continuous on its domain, convex in $\phi \in \mathcal{F}$, concave in $(\mu, \nu) \in \mathcal{M}_1 \times \mathcal{M}_2$ and possesses saddle point (min in ϕ , max in (μ, ν)):

there exist $(\phi_* \in \mathcal{F}, (\mu^*, \nu^*) \in \mathcal{M}_1 \times \mathcal{M}_2)$ such that

$$\Phi(\phi; \mu^*, \nu^*) \geq \underbrace{\Phi(\phi_*; \mu^*, \nu^*)}_{=\text{SadVal}} \geq \Phi(\phi_*; \mu, \nu)$$

for all $(\phi \in \mathcal{F}, (\mu, \nu) \in \mathcal{M}_1 \times \mathcal{M}_2)$.

- The component ϕ_* of a saddle point $(\phi_*, (\mu^*, \nu^*))$ of Φ is an optimal solution to (!), and

$$\varepsilon_*(\mathcal{P}_1, \mathcal{P}_2) = \exp\{\Phi(\phi_*; \mu^*, \nu^*)\}.$$

- A saddle point $(\phi_*, (\mu^*, \nu^*))$ can be found as follows. We solve the optimization problem

$$\text{SadVal} = \max_{\mu \in M_1, \nu \in M_2} \ln \left(\int_{\Omega} \sqrt{p_{\mu}(\omega) p_{\nu}(\omega)} \Pi(d\omega) \right)$$

(which is a solvable convex optimization problem), and take an optimal solution to the problem as (μ^*, ν^*) . We then set

$$\phi_*(\omega) = \frac{1}{2} \ln \left(p_{\mu^*}(\omega) / p_{\nu^*}(\omega) \right),$$

thus getting an optimal detector $\phi_* \in \mathcal{F}$.

- For this detector and the associated simple test $\mathcal{T}_{\phi_*}$,

$$\begin{aligned} \text{Risk}[\mathcal{T}_{\phi_*} | H_1, H_2] &\leq \text{Risk}[\phi_* | \mathcal{P}_1, \mathcal{P}_2] \\ &= \text{Risk}_1[\phi_* | \mathcal{P}_1, \mathcal{P}_2] = \text{Risk}_2[\phi_* | \mathcal{P}_1, \mathcal{P}_2] \\ &= \varepsilon_*(\mathcal{P}_1, \mathcal{P}_2) = e^{\text{SadVal}} = \int_{\Omega} \sqrt{p_{\mu^*}(\omega) p_{\nu^*}(\omega)} \Pi(d\omega). \end{aligned}$$

$$\text{SadVal} = \max_{\mu \in M_1, \nu \in M_2} \ln \left(\int_{\Omega} \sqrt{p_{\mu}(\omega) p_{\nu}(\omega)} \Pi(d\omega) \right) \quad (!)$$

Gaussian o.s. $\mathcal{P}_{\chi} = \{\mathcal{N}(\mu, I_d) : \mu \in M_{\chi}\}$, $\chi = 1, 2$:

- Optimization problem (!) reads

$$\max_{\mu \in M_1, \nu \in M_2} \left[-\frac{1}{8} \|\mu - \nu\|_2^2 \right]$$

- The optimal balanced detector and its risk are given by

$$\left| \begin{aligned} \phi_*(\omega) &= \frac{1}{2}[\mu^* - \nu^*]\omega - c, \\ \left[(\mu^*, \nu^*) \in \underset{\mu \in M_1, \nu \in M_2}{\text{Argmin}} \|\mu - \nu\|_2^2, \quad c &= \frac{1}{4}[\mu^* - \nu^*]^T [\mu^* + \nu^*] \right] \\ \varepsilon_*(\mathcal{P}_1, \mathcal{P}_2) &= \exp \left\{ -\frac{\|\mu^* - \nu^*\|_2^2}{8} \right\} \end{aligned} \right|$$

Note: We are in the “signal plus noise” model of observations with noise $\sim \mathcal{N}(0, I_d)$. The test $\mathcal{T}_{\phi_*}$ is nothing but the pairwise Euclidean separation test associated with $X_{\chi} = M_{\chi}$, $\chi = 1, 2$.

$$\text{SadVal} = \max_{\mu \in M_1, \nu \in M_2} \ln \left(\int_{\Omega} \sqrt{p_{\mu}(\omega)p_{\nu}(\omega)} \Pi(d\omega) \right) \quad (!)$$

Poisson o.s. $\mathcal{P}_{\chi} = \left\{ \bigotimes_{i=1}^d \text{Poisson}(\mu_i) : \mu = [\mu_1; \dots; \mu_d] \in M_{\chi} \right\}, \chi = 1, 2:$

- Problem (!) reads

$$\max_{\mu \in M_1, \nu \in M_2} \underbrace{\left[-\frac{1}{2} \sum_{i=1}^d (\sqrt{\mu_i} - \sqrt{\nu_i})^2 \right]}_{\sum_i [\sqrt{\mu_i \nu_i} - \frac{1}{2} \mu_i - \frac{1}{2} \nu_i]}$$

- The optimal balanced detector and its risk are given by

$$\left| \begin{array}{l} \phi_*(\omega) = \frac{1}{2} \sum_{i=1}^d [\ln(\mu_i^*/\nu_i^*) \omega_i + \nu_i^* - \mu_i^*], \\ \left[(\mu^*, \nu^*) \in \underset{\mu \in M_1, \nu \in M_2}{\text{Argmax}} \sum_i [\sqrt{\mu_i \nu_i} - \frac{1}{2} \mu_i - \frac{1}{2} \nu_i] \right] \\ \varepsilon_*(\mathcal{P}_1, \mathcal{P}_2) = \exp \left\{ -\frac{1}{2} \sum_i (\sqrt{\mu_i^*} - \sqrt{\nu_i^*})^2 \right\} \end{array} \right|$$

$$\text{SadVal} = \max_{\mu \in M_1, \nu \in M_2} \ln \left(\int_{\Omega} \sqrt{p_{\mu}(\omega) p_{\nu}(\omega)} \Pi(d\omega) \right) \quad (!)$$

Discrete o.s.

$$\mathcal{P}_{\chi} = \{\mu \in M_{\chi}\}, M_{\chi} \subset \Delta_d^o = \{\mu \in \mathbb{R}_+^d : \sum_{\omega} \mu_{\omega} = 1, \mu > 0\}, \chi = 1, 2$$

- Problem (!) reads

$$\max_{\mu \in M_1, \nu \in M_2} \sum_{\omega} \sqrt{\mu_{\omega} \nu_{\omega}}$$

- The optimal balanced detector and its risk are given by

$$\left| \begin{array}{l} \phi_*(\omega) = \frac{1}{2} \ln(\mu_{\omega}^* / \nu_{\omega}^*), \quad \omega \in \Omega = \{1, \dots, d\} \\ \left[(\mu^*, \nu^*) \in \underset{\mu \in M_1, \nu \in M_2}{\text{Argmin}} \sum_{\omega} \sqrt{\mu_{\omega} \nu_{\omega}} \right] \\ \varepsilon_*(\mathcal{P}_1, \mathcal{P}_2) = \sum_{\omega} \sqrt{\mu_{\omega}^* \nu_{\omega}^*} \end{array} \right.$$

Direct product of simple o.s.'s. Let

$$\mathcal{O}_k = \left((\Omega_k, \Pi_k), \{p_{\mu_k, k}(\cdot) : \mu_k \in \mathcal{M}_k\}, \mathcal{F}_k \right), \quad 1 \leq k \leq K,$$

be simple o.s.'s, and $M_{\chi, k} \subset \mathcal{M}_k$, $\chi = 1, 2$, be nonempty convex compact sets. Consider the simple o.s.

$$\left((\Omega^K, \Pi^K), \{p_{\mu^K} : \mu^K \in \mathcal{M}^K\}, \mathcal{F}^K \right) = \bigotimes_{k=1}^K \mathcal{O}_k$$

along with two compact convex sets

$$M_\chi = M_{\chi, 1} \times \dots \times M_{\chi, K}, \quad \chi = 1, 2.$$

Question: *What is the problem*

$$\max_{\mu^K \in M_1, \nu^K \in M_2} \ln \left(\int_{\Omega^K} \sqrt{p_{\mu^K}(\omega^K) p_{\nu^K}(\omega^K)} \Pi^K(d\omega^K) \right)$$

responsible for the smallest risk detector for the families of distributions $\mathcal{P}_\chi^K$, $\mathcal{P}_2^K$ associated in $\mathcal{O}^K$ with the sets M_χ , $\chi = 1, 2$?

Answer: *This is the separable problem*

$$\max_{\{\mu_k \in M_{1,k}, \nu_k \in M_{2,k}\}_{k=1}^K} \sum_{k=1}^K \ln \left(\int_{\Omega_k} \sqrt{p_{\mu_k,k}(\omega_k) p_{\nu_k,k}(\omega_k)} \Pi_k(d\omega_k) \right)$$

$\Rightarrow$ *Minimum risk balanced detector for $\mathcal{P}_1^K, \mathcal{P}_2^K$ can be chosen as*

$$\begin{aligned} \phi_*^K(\omega_1, \dots, \omega_K) &= \sum_{k=1}^K \phi_{*,k}(\omega_k), \\ \phi_{*,k}(\omega_k) &= \frac{1}{2} \ln \left(p_{\mu_k^*,k}(\omega) / p_{\nu_k^*,k}(\omega) \right) \end{aligned}$$

$$\left[(\mu_k^*, \nu_k^*) \in \underset{\mu_k \in M_{1,k}, \nu_k \in M_{2,k}}{\text{Argmax}} \ln \left(\int_{\Omega_k} \sqrt{p_{\mu_k,k}(\omega_k) p_{\nu_k,k}(\omega_k)} \Pi_k(d\omega_k) \right) \right]$$

and

$$\varepsilon_*(\mathcal{P}_1^K, \mathcal{P}_2^K) = \prod_{k=1}^K \varepsilon_*(\mathcal{P}_{1,k}, \mathcal{P}_{2,k}),$$

where $\mathcal{P}_{\chi,k}$ are the families of distributions associated in $\mathcal{O}_k$ with $M_{\chi,k}$, $\chi = 1, 2$.

Power of a simple o.s. Let

$$\mathcal{O} = ((\Omega, \Pi), \{p_\mu(\cdot) : \mu \in \mathcal{M}\}, \mathcal{F})$$

be a simple o.s., and $M_\chi \subset \mathcal{M}$, $\chi = 1, 2$, be nonempty convex compact sets. Consider the K -th power of $\mathcal{O}$, that is, the simple o.s.

$$\mathcal{O}^{(K)} = \left((\Omega^K, \Pi^K), \{p_\mu^{(K)}(\omega_1, \dots, \omega_K) = \prod_{k=1}^K p_\mu(\omega_k) : \mu \in \mathcal{M}\}, \mathcal{F}^{(K)} \right).$$

Question: *What is the problem*

$$\max_{\mu \in M_1, \nu \in M_2} \ln \left(\int_{\Omega^K} \sqrt{p_\mu^{(K)}(\omega^K) p_\nu^{(K)}(\omega^K)} \Pi^K(d\omega^K) \right)$$

responsible for the smallest risk detector for the families of distributions $\mathcal{P}_\chi^{(K)}$ associated in $\mathcal{O}^{(K)}$ with the sets M_χ , $\chi = 1, 2$?

Answer: *This is the separable problem*

$$\max_{\{\mu \in M_1, \nu \in M_2\}} \underbrace{\sum_{k=1}^K \ln \left(\int_{\Omega} \sqrt{p_{\mu}(\omega_k) p_{\nu}(\omega_k)} \Pi(d\omega_k) \right)}_{K \ln \left(\int_{\Omega} \sqrt{p_{\mu}(\omega) p_{\nu}(\omega)} \Pi(d\omega) \right)}$$

$\Rightarrow$ *Minimum risk balanced detector for $\mathcal{P}_1^{(K)}, \mathcal{P}_2^{(K)}$ can be chosen as*

$$\begin{aligned} \phi_*^{(K)}(\omega_1, \dots, \omega_K) &= \sum_{k=1}^K \phi_*(\omega_k), \\ \phi_*(\omega_k) &= \frac{1}{2} \ln \left(p_{\mu^*}(\omega) / p_{\nu^*}(\omega) \right) \\ \left[(\mu^*, \nu^*) \in \underset{\mu \in M_1, \nu \in M_2}{\text{Argmax}} \ln \left(\int_{\Omega} \sqrt{p_{\mu}(\omega) p_{\nu}(\omega)} \Pi(d\omega) \right) \right] \end{aligned}$$

and

$$\varepsilon_*(\mathcal{P}_1^{(K)}, \mathcal{P}_2^{(K)}) = [\varepsilon_*(\mathcal{P}_1, \mathcal{P}_2)]^K,$$

where $\mathcal{P}_{\chi}$ are the families of distributions associated in $\mathcal{O}$ with M_{χ} , $\chi = 1, 2$.

Near-Optimality of Detector-Based Tests in Simple O.S.

Proposition A. *Let*

$$\mathcal{O} = ((\Omega, \Pi), \{p_\mu : \mu \in \mathcal{M}\}, \mathcal{F})$$

be a simple o.s., $M_\chi \subset \mathcal{M}$, $\chi = 1, 2$, be nonempty convex compact sets, giving rise to families of distributions

$$\mathcal{P}_\chi = \{P : P \text{ has density } p_\mu(\cdot) \text{ w.r.t. } \Pi \text{ with } \mu \in M_\chi\}, \chi = 1, 2,$$

hypotheses $H_\chi : P \in \mathcal{P}_\chi$, $\chi = 1, 2$, on the distribution of a random observation $\omega \in \Omega$, and minimum risk detector ϕ_ for $\mathcal{P}_1, \mathcal{P}_2$.*

Assume that in the nature there exists a simple single-observation test, deterministic or randomized, $\mathcal{T}$ with

$$\text{Risk}[\mathcal{T}|H_1, H_2] \leq \epsilon < 1/2.$$

Then the risk of the simple test $\mathcal{T}_{\phi_}$ accepting H_1 when $\phi_*(\omega) \geq 0$ and accepting H_2 otherwise “is comparable” to ϵ :*

$$\text{Risk}[\mathcal{T}_{\phi_*}|H_1, H_2] \leq \bar{\epsilon} := 2\sqrt{\epsilon(1-\epsilon)} < 1.$$

Proof. From what we called “universality” of detector-based tests, there exists a detector ϕ with $\text{Risk}[\phi|\mathcal{P}_1, \mathcal{P}_2] \leq \bar{\epsilon}$, and $\text{Risk}[\phi_*|\mathcal{P}_1, \mathcal{P}_2]$ can be only less than $\text{Risk}[\phi|\mathcal{P}_1, \mathcal{P}_2]$. $\square$

Proposition B. Let $\mathcal{O} = ((\Omega, \Pi), \{p_\mu : \mu \in \mathcal{M}\}, \mathcal{F})$ be a simple o.s., and $M_\chi \subset \mathcal{M}$, $\chi = 1, 2$, be nonempty convex compact sets, giving rise to families of distributions

$$\mathcal{P}_\chi = \{P : P \text{ has density } p_\mu(\cdot) \text{ w.r.t. } \Pi \text{ with } \mu \in M_\chi\}, \chi = 1, 2$$

their powers

$$\mathcal{P}_\chi^K = \{P \times \dots \times P : P \in \mathcal{P}_\chi\}, \chi = 1, 2, K = 1, 2, \dots$$

hypotheses

$$H_\chi^K : P \in \mathcal{P}_\chi^K, \chi = 1, 2, K = 1, 2, \dots$$

on the distribution of a random K -repeated observation $\omega^K = (\omega_1, \dots, \omega_K) \in \Omega^K$, and minimum risk detector ϕ_* for $\mathcal{P}_1, \mathcal{P}_2$.

Assume that “in the nature” there exists positive integer K_* and a simple K_* -observation test, deterministic or randomized, $\mathcal{T}_{K_*}$ capable to decide on the hypotheses $H_\chi^{K_*}$, $\chi = 1, 2$, with risk $\leq \epsilon < 1/2$. Then the test $\mathcal{T}_{\phi_*, K}$ deciding on H_χ^K , $\chi = 1, 2$, by accepting H_1^K whenever

$$\phi^{(K)}(\omega^K) := \sum_{k=1}^K \phi_*(\omega_k) \geq 0$$

and accepting H_2^K otherwise, satisfies

$$\text{Risk}[\mathcal{T}_{\phi_*}, |H_1^K, H_2^K] \leq \epsilon \quad \forall K \geq \hat{K}_* = \frac{2}{1 - \frac{\ln(4(1-\epsilon))}{\ln(1/\epsilon)}} K_*. \quad \left(\hat{K}_*/K_* \rightarrow 2 \text{ as } \epsilon \rightarrow +0 \right).$$

Moreover, this risk bound remains true when the hypotheses H_χ^K are extended to hypotheses $H_\chi^{\oplus, K}$ stating that the distribution P of ω^K belongs to the quasi direct K -th power of $\mathcal{P}_\chi$, $\chi = 1, 2$.

Proof. As we know, K_* -th power $\mathcal{O}^{(K)}$ of $\mathcal{O}$ is simple o.s. along with $\mathcal{O}$, and $\phi_*^{(K)}$ is the minimum risk detector for the families $\mathcal{P}_\chi^{K_*}$, $\chi = 1, 2$, the risk of this detector being $[\varepsilon_*(\mathcal{P}_1, \mathcal{P}_2)]^{K_*}$. By Proposition A as applied to $\mathcal{O}^{(K)}$ in the role of $\mathcal{O}$, we have

$$[\varepsilon_*(\mathcal{P}_1, \mathcal{P}_2)]^{K_*} \leq 2\sqrt{\epsilon(1-\epsilon)} \Rightarrow \varepsilon_*(\mathcal{P}_1, \mathcal{P}_2) \leq [2\sqrt{\epsilon(1-\epsilon)}]^{1/K_*} < 1.$$

By Detector Calculus, it follows that for $K = 1, 2, \dots$ it holds

$$\text{Risk}[\phi_*^{(K)} | \mathcal{P}_1^{\oplus, K}, \mathcal{P}_2^{\oplus, K}] \leq [2\sqrt{\epsilon(1-\epsilon)}]^{K/K_*}$$

and the right hand side is $\leq \epsilon$ whenever $K \geq \widehat{K}_*$. □

Near-Optimality of Detector-Based Up to Closeness Testing

Situation: We are given a simple o.s.

$$\mathcal{O} = ((\Omega, \Pi), \{p_\mu : \mu \in \mathcal{M}\}, \mathcal{F})$$

and a collection of nonempty convex compact subsets M_ℓ , $1 \leq \ell \leq L$ giving rise to

- Families

$$\mathcal{P}_\ell = \{P : P \text{ admits density } p_\mu, \mu \in M_\ell \text{ w.r.t. } \Pi\}, \ell = 1, \dots, L,$$

along with direct products $\mathcal{P}_\ell^{\oplus, K} = \bigoplus_{k=1}^K \mathcal{P}_\ell$ of $\mathcal{P}_\ell$ and hypotheses $H_\ell^{\oplus, K} : P^K \in \mathcal{P}_\ell^{\oplus, K}$ on the distribution of K -repeated observation $\omega^K = (\omega_1, \dots, \omega_K)$,

- minimum-risk balanced single-observation detectors $\phi_{\ell\ell'}(\omega)$ for $\mathcal{P}_\ell, \mathcal{P}_{\ell'}$ along with their risks $\varepsilon_*(\mathcal{P}_\ell, \mathcal{P}_{\ell'})$, $1 \leq \ell < \ell' \leq L$, and K -repeated versions

$$\phi_{\ell\ell'}^K(\omega^K) = \sum_{k=1}^K \phi_{\ell\ell'}(\omega_k)$$

of $\phi_{\ell\ell'}$ such that $\text{Risk}[\phi_{\ell\ell'}^{(K)} | H_\ell^{\oplus, K}, H_{\ell'}^{\oplus, K}] \leq [\varepsilon_*(\mathcal{P}_\ell, \mathcal{P}_{\ell'})]^K$.

- Assume that in addition to the above data, we are given a closeness relation $\mathcal{C}$ on $\{1, \dots, L\}$. Applying Calculus of Detectors, for every positive integer K , setting

$$\theta_K = \left\| \left[\varepsilon_*^K(\mathcal{P}_\ell, \mathcal{P}_{\ell'}) \cdot \begin{cases} 1, & (\ell, \ell') \notin \mathcal{C} \\ 0, & (\ell, \ell') \in \mathcal{C} \end{cases} \right]_{\ell, \ell'=1}^L \right\|_{2,2}$$

we can assemble the outlined data, in a computationally efficient fashion, into a K -observation test $\mathcal{T}^K$ deciding on $H_\ell^{\oplus, K}$, $1 \leq \ell \leq L$, with $\mathcal{C}$ -risk upper-bounded as follows:

$$\text{Risk}^{\mathcal{C}}[\mathcal{T}^K | H_1^{\oplus, K}, \dots, H_L^{\oplus, K}] \leq \varkappa \theta_K$$

($\varkappa > 1$ can be selected to be as close to 1 as we want).

Proposition. *In the just described situation, assume that for some $\epsilon < 1/2$ and K_* in the nature there exists a K_* -observation test $\mathcal{T}$, deterministic or randomized, deciding on the hypotheses*

$H_\ell^{K_} : \omega^{K_*} = (\omega_1, \dots, \omega_{K_*})$ is an i.i.d. sample drawn from a $P \in \mathcal{P}_\ell$, $\ell = 1, \dots, L$, with $\mathcal{C}$ -risk $\leq \epsilon$. Then the test $\mathcal{T}^K$ with*

$$K \geq 2 \underbrace{\left[\frac{1 + \ln(\kappa L) / \ln(1/\epsilon)}{1 - \ln(4(1 - \epsilon)) / \ln(1/\epsilon)} \right]}_{\rightarrow 1 \text{ as } \epsilon \rightarrow +0} K_*$$

decides on $H_\ell^{\oplus, K}$, $\ell = 1, \dots, L$, with $\mathcal{C}$ -risk $\leq \epsilon$ as well.

Proof. 1^o Let us fix ℓ, ℓ' such that $(\ell, \ell') \notin \mathcal{C}$, and let us convert $\mathcal{T}$ into a simple K_* -observation test $\tilde{\mathcal{T}}_{\ell, \ell'}$ deciding on $H_\ell^{K_*}, H_{\ell'}^{K_*}$ as follows: whenever $\ell \in \mathcal{T}(\omega^K)$, $\tilde{\mathcal{T}}_{\ell, \ell'}$ accepts $H_\ell^{K_*}$ and rejects $H_{\ell'}^{K_*}$, otherwise the test accepts $H_{\ell'}^{K_*}$ and rejects $H_\ell^{K_*}$. It is immediately seen that

$$\text{Risk}[\tilde{\mathcal{T}}_{\ell, \ell'} | H_\ell^{K_*}, H_{\ell'}^{K_*}] \leq \epsilon.$$

Indeed, let $P^{K_*} = P \times \dots \times P$ be the distribution of ω^{K_*} . Whenever P^K obeys $H_\ell^{K_*}$, $\mathcal{T}$ must accept the hypothesis with P^{K_*} -probability $\geq 1 - \epsilon$, whence

$$\text{Risk}_1[\tilde{\mathcal{T}}_{\ell, \ell'} | H_\ell^{K_*}, H_{\ell'}^{K_*}] \leq \epsilon.$$

If P^{K_*} obeys $H_{\ell'}^{K_*}$, the P^{K_*} -probability of the event “ $\mathcal{T}$ accepts $H_{\ell'}^{K_*}$ and rejects $H_\ell^{K_*}$ ” is $\leq \epsilon$, since $H_{\ell'}^{K_*}, H_\ell^{K_*}$ are not $\mathcal{C}$ -close to each other

$\Rightarrow$ P^{K_*} -probability to reject $H_\ell^{K_*}$ is at least $1 - \epsilon$

$\Rightarrow \text{Risk}_2[\tilde{\mathcal{T}}_{\ell, \ell'} | H_\ell^{K_*}, H_{\ell'}^{K_*}] \leq \epsilon.$

2° Because $H_{\ell'}^{K*}, H_{\ell}^{K*}$ can be decided upon with risk $\leq \epsilon < 1/2$, by detector calculus

$$\varepsilon_*(\mathcal{P}_{\ell}^{K*}, \mathcal{P}_{\ell'}^{K*}) \leq 2\sqrt{\epsilon(1-\epsilon)} < 1$$

Since $\mathcal{O}$ is a simple o.s,

$$\varepsilon_*(\mathcal{P}_{\ell}, \mathcal{P}_{\ell'}) \leq \left[2\sqrt{\epsilon(1-\epsilon)}\right]^{1/K*} < 1,$$

We conclude that

$$\theta_K \leq \left[2\sqrt{\epsilon(1-\epsilon)}\right]^{K/K*} L$$

and

$$\text{Risk}^{\mathcal{C}}[\mathcal{T}^K | H_1^{\oplus, K}, \dots, H_L^{\oplus, K}] \leq \varkappa \theta_K \leq \epsilon$$

when

$$\frac{K}{K*} \geq 2 \frac{1 + \ln(\varkappa L) / \ln(1/\epsilon)}{1 - \ln(4(1-\epsilon)) / \ln(1/\epsilon)}.$$

□

Illustration: Selecting the Best in a Family of Estimates

Problem: we are given

- a simple o.s. $\mathcal{O} = ((\Omega, \pi), \{p_\mu : \mu \in \mathcal{M}\}, \mathcal{F})$ and have access to stationary K -repeated observations

$$\omega_k \sim p_{A(x_*)}(\cdot), k = 1, \dots, K,$$

of unknown signal x_* known to belong to a given convex compact set $X \subset \mathbb{R}^n$.

Here $[x \mapsto A(x)]$: affine mapping such that $A(X) \subset \mathcal{M}$;

- M candidate estimates $x_i \in \mathbb{R}^n$, $1 \leq i \leq M$, of x_* ,
- a norm $\|\cdot\|$ on $\mathbb{R}^n$, and a reliability tolerance $\epsilon \in (0, 1)$

Parameter space

Given observations $\omega_1, \dots, \omega_K$ we look to identify $(1 - \epsilon)$ -reliably the $\| \cdot \|$ -closest to x_ point among $x_1, \dots, x_M$.*

Actual goal: given $\alpha \geq 1$, $\beta \geq 0$ and a grid

$$\Gamma = \{r_0 > r_1 > \dots > r_N > 0\},$$

use observations $\omega_1, \dots, \omega_K$ to identify $(1 - \epsilon)$ -reliably $x_{i(\omega^K)}$ such that

$$\|x_* - x_{i(\omega^K)}\| \leq \alpha \rho(x_*) + \beta$$

where

$$\rho(x) := \min\{r : r \in \Gamma, r \geq \min_i \|x - x_i\|\}$$

is the “grid approximation” of $\min_i \|x - x_i\|$.

Parameter space

Observation parameter space

Solution:

- Consider $M(N + 1)$ hypotheses

$$H_{ij} : \omega_k \sim p_{A(x)}(\cdot) \text{ for some } x \in X_{ij} := \{x \in X : \|x - x_i\| \leq r_j\}.$$

- Define closeness $\mathcal{C} = \mathcal{C}_{\alpha, \beta}$:

$$[ij] \text{ is } \mathcal{C}\text{-close to } [i'j'] \text{ iff } \|x_i - x_{i'}\| \leq r_j + r_{j'} + \beta.$$

- Build minimum risk pairwise detectors $\phi_{ij, i'j'}$, along with their risks $\epsilon_{ij, i'j'}$, for hypotheses H_{ij} .
- Choose β such that $\epsilon_{ij, i'j'} < 1$ for all $(ij, i'j') \notin \mathcal{C}$.
- Find the smallest $K = K_\epsilon(\beta)$ such that the K -observation detector-based test $\mathcal{T}^K$ decides on $\{H_{ij} : i \leq M, 0 \leq j \leq N\}$ with $\mathcal{C}$ -risk $\leq \epsilon$.
- To build $i(\omega^K)$,
 - apply $\mathcal{T}^K$ to observation ω^K ,
 - select $H_{i_*(\omega^K)j_*(\omega^K)}$ with the largest j among accepted hypotheses H_{ij} .

Proposition. *In the situation in question, assume that for some $\epsilon \in (0, 1/2)$, $b \geq 0$ and positive integer $\bar{K}$ in the nature there exists an inference $\omega^{\bar{K}} \rightarrow i_*(\omega^{\bar{K}})$ such that, for all $x^* \in X$,*

$$\text{Prob}\{\|x_* - x_{i_*(\omega^{\bar{K}})}\| \leq \rho(x_*) + b\} \geq 1 - \epsilon.$$

- *Then one can choose $\beta = 2b$.*
- *The resulting inference $\omega^K \mapsto i(\omega^K)$ satisfies*

$$\text{Prob}\{\|x_* - x_{i(\omega^K)}\| \leq 3\rho(x_*) + 2b\} \geq 1 - \epsilon.$$

for $K \geq K_\epsilon(2b)$, where

$$K_\epsilon(2b) \leq \text{Ceil} \left(\frac{2 \ln(M(N+1))/\ln(1/\epsilon)}{1 - \ln(4(1-\epsilon))/\ln(1/\epsilon)} \bar{K} \right).$$

Numerical illustration. Given noisy observation

$$\omega = Ax + \sigma\xi, \quad \xi \sim \mathcal{N}(0, I_n)$$

of the “discretized primitive” Ax of a signal $x = [x^1; \dots; x^n] \in \mathbb{R}^n$:

$$[Ax]_j = \frac{1}{n} \sum_{s=1}^j x^s, \quad 1 \leq j \leq n,$$

for $i = 1, \dots, \kappa$ we have built Least Squares polynomial, of order $i - 1$, approximations x_i of x :

$$x_i = \operatorname{argmin}_{x \in \mathcal{P}_i} \|Ax - \omega\|_2^2$$
$$\left[\mathcal{P}_i = \{x = [x^1; \dots; x^n] : x^s \text{ is polynomial, of degree } \leq i - 1, \text{ in } s\} \right]$$

and now want to use K additional observations to identify the nearly closest to x_* , in the norm

$$\|u\| = \frac{1}{n} \sum_{i=1}^n |u_i|$$

on $\mathbb{R}^n$, among the points x_i , $1 \leq i \leq \kappa$.

Numerical experiment [$\epsilon = 0.01, n = 128, \sigma = 0.01, \kappa = 5, \beta = 0.05$]

Top: x_* and x_i . Bottom: the primitive of x_*

i	1	2	3	4	5
$\ x - x_i\ $	0.5348	0.33947	0.23342	0.16313	0.16885

distances from x_* to x_i

- With $K = 3$ (computed value), for a sample of size 1000, *the closest to x_* point was correctly identified in all simulations*
- the same was true when $K = 3$ was replaced with $K = 1$;
- replacing $K = 3$ with $K = 1$ *and increasing σ from 0.01 to 0.05*, the procedure started to make imperfect conclusions.

However, the exactly closest to x_* point x_4 was identified correctly in as many as 96% of simulations, with estimated **excess error**

$$\mathbf{E}\{\|x_* - x_{i(\omega^1)}\| - \rho(x_*)\} \leq 0.0024.$$

Illustration: Recovering Linear-Fractional Function of a Signal

Problem: An unknown signal x known to belong to a given convex compact set $X \subset \mathbb{R}^n$ is observed according to

$$\omega = Ax + \sigma\xi, \quad \xi \sim \mathcal{N}(0, I_d)$$

Our goal is to recover the value at x of a linear-fractional functional

$$F(z) = \frac{f^T z}{e^T z}, \quad \text{with } e^T z > 0, z \in X.$$

Illustration: We are given noisy measurements of voltages V_i at *some* nodes i and currents I_{ij} in *some* arcs (i, j) of an electric circuit, and want to recover the resistance of a particular arc $(\hat{i}, \hat{j})$:

$$r_{\hat{i}\hat{j}} = \frac{V_{\hat{j}} - V_{\hat{i}}}{I_{\hat{i}\hat{j}}}$$

Circuit with 8 nodes and 11 arcs

$$x = [\text{voltages at nodes; currents in arcs}]$$

$$Ax = [\text{observable voltages; observable currents}]$$

- Currents are measured in blue arcs only
- Voltages are measured in magenta nodes only
- We want to recover resistance of red arc
- $X : \left\{ \begin{array}{l} \text{conservation of current, except for nodes \#1,8} \\ \text{zero voltage at node \#1, nonnegative currents} \\ \text{current in red arc at least 1, total of currents at most 33} \\ \text{Ohm Law, resistances of arcs between 1 and 10} \end{array} \right.$

Strategy: Given L ,

- split the range

$$\Delta = [\min_{x \in X} F(x), \max_{x \in X} F(x)]$$

into L consecutive bins Δ_ℓ of length $\delta_L = \text{length}(\Delta)/L$,

- define convex compact sets

$$X_\ell = \{x \in X : F(x) \in \Delta_\ell\}, \quad M_\ell = \{Ax : x \in X_\ell\}, \quad 1 \leq \ell \leq L$$

a 2D projection of X and $X_1, \dots, X_8$

- We are to decide on L hypotheses

$$H_\ell : P = \mathcal{N}(\mu, \sigma^2 I), \quad \mu \in M_\ell$$

about the distribution P of observation

$$\omega = Ax + \sigma\xi$$

up to closeness

$$\mathcal{C} : \text{“} H_\ell \text{ is close to } H_{\ell'} \text{ if and only if } |\ell - \ell'| \leq 1 \text{”}$$

- Estimate $F(x)$ by the barycenter of accepted bins.

Fact: *For the resulting test $\mathcal{T}$, with probability*

$$\geq 1 - \text{Risk}^{\mathcal{C}}[\mathcal{T} | H_1, \dots, H_L]$$

the estimation error does not exceed δ_L .

Implementation and results: Given target risk ϵ and L , we select the largest σ for which $\text{Risk}^{\mathcal{C}}[\mathcal{T}|H_1, \dots, H_L]$ is $\leq \epsilon$.

- This is what we get for $\epsilon = 0.01$: $\Delta = [1, 10]$

L	8	16	32
δ_L	$9/8 \approx 1.13$	$9/16 \approx 0.56$	$9/32 \approx 0.28$
σ	0.024	0.010	0.005
$\sigma_{\text{opt}}/\sigma \leq$	1.31	1.31	1.33
σ	0.031	0.013	0.006
$\sigma_{\text{opt}}/\sigma \leq$	1.01	1.06	1.08

- σ_{opt} – the largest σ for which “in the nature” there exists a test deciding on $H_1, \dots, H_L$ with $\mathcal{C}$ -risk ≤ 0.01

- **Red data:** Risks $\epsilon_{\ell\ell'}$ of pairwise tests are bounded via risks of optimal detectors, $\mathcal{C}$ -risk of $\mathcal{T}$ is bounded by

$$\left\| \left[\epsilon_{\ell\ell'} \cdot \chi_{(\ell,\ell') \notin \mathcal{C}} \right]_{\ell,\ell'=1}^L \right\|_{2,2};$$

- **Black data:** Risks $\epsilon_{\ell\ell'}$ of pairwise tests are bounded via error function, $\mathcal{C}$ -risk of $\mathcal{T}$ is bounded by

$$\max_{\ell} \sum_{\ell': (\ell,\ell') \notin \mathcal{C}} \epsilon_{\ell\ell'}.$$

Recovering N -Convex Functionals on Unions of Convex Sets

Fact: The construction used to recover linear-fractional function can be extended to recovering N -convex functionals.

Definition: Let $X \subset \mathbb{R}^n$ be a convex compact set, $F : X \rightarrow \mathbb{R}$ be a continuous function, and N be a positive integer. We say that F is N -convex, if for every real a the sets

$$X^{a,\geq} = \{x \in X : F(x) \geq a\}, \quad X^{a,\leq} = \{x \in X : F(x) \leq a\}$$

can be represented as the unions of at most N convex compact sets.

Example A: Fractional-linear function $F(x) = \frac{e(x)}{d(x)}$ with positive on X denominator is 1-convex:

$$\{x \in X : F(x) \underset{\leq}{\overset{\geq}{\geq}} a\} = \{x \in X : e(x) - ad(x) \underset{\leq}{\overset{\geq}{\geq}} 0\}$$

Example B: If F_i is N_i -convex on X , $i = 1, 2$, then $\max[F_1, F_2]$ and $\min[F_1, F_2]$ are $\max[N_1 + N_2, N_1 N_2]$ -convex on X :

$$\begin{cases} X_i^{a, \geq} := \{x \in X : F_i(x) \geq a\} = \bigcup_{\nu=1}^{N_i} U_{\nu, i}^a \\ X_i^{a, \leq} := \{x \in X : F_i(x) \leq a\} = \bigcup_{\nu=1}^{N_i} V_{\nu, i}^a \end{cases}, i = 1, 2 \text{ } [U, V: \text{convex}]$$

$$\Rightarrow \begin{cases} \{x \in X : \max[F_1(x), F_2(x)] \geq a\} = \left[\bigcup_{\mu \leq N_1} U_{\mu, 1}^a \right] \cup \left[\bigcup_{\nu \leq N_2} U_{\nu, 2}^a \right] \\ \{x \in X : \max[F_1(x), F_2(x)] \leq a\} = \bigcup_{\mu \leq N_1, \nu \leq N_2} \left[V_{\mu, 1}^a \cap V_{\nu, 2}^a \right] \end{cases}$$

Extension: If F_i is N_i -convex on X , $i = 1, 2, \dots, M$, then $\max[F_1, F_2, \dots, F_M]$ and $\min[F_1, F_2, \dots, F_M]$ are $\max[\sum_i N_i, \prod_i N_i]$ -convex on X .

Problem of interest: *Given*

- *convex compact set $X \subset \mathbb{R}^n$,*
- *N -convex functional $F : X \rightarrow \mathbb{R}$,*
- *a collection $X_j, j = 1, \dots, J$, of convex compact subsets of X ,*
- *stationary K -repeated observations $\omega_1, \dots, \omega_K$ stemming, via simple o.s.,*
from unknown signal $x \in \bigcup_{j=1}^J X_j$,

we want to recover $F(x)$.

Strategy: Given L , we

- split the range

$$\Delta = [\min_{x \in X} F(x), \max_{x \in X} F(x)]$$

into L consecutive bins Δ_ℓ of length $\delta_L = \text{length}(\Delta)/L$,

- and observe that by N -convexity of F , each set

$$\{x \in \bigcup_{j=1}^J X_j : F(x) \in \Delta_\ell\}$$

is the union of at most $N^2 J$ convex compact sets Y_s^ℓ ;

- Associate with nonempty sets Y_s^ℓ the hypotheses “**observation stems from a signal from Y_s^ℓ** ”
- Define closeness $\mathcal{C}$ on the resulting collection of hypotheses $H_1, \dots, H_{\mathcal{L}}$, $\mathcal{L} \leq N^2 J L$:

H_μ and H_ν are $\mathcal{C}$ -close

iff both hypotheses stem from the same or from two consecutive bins Δ_ℓ

- Use our gear for testing multiple convex hypotheses in simple o.s. to build a test $\mathcal{T}_K$ deciding on $H_1, \dots, H_{\mathcal{L}}$ up to closeness $\mathcal{C}$ via K -repeated observation.
- We apply the test $\mathcal{T}_K$ to observations $\omega_1, \dots, \omega_K$ and take as the estimate of $F(x)$ the center of masses of all bins associated with the hypotheses accepted by the test.

Same as in the above fractional-linear example, it is immediately seen that

The probability for the recovery error to be $> \delta_L$ is at most the $\mathcal{C}$ -risk of $\mathcal{T}_K$.

In addition, *with our estimate, the number of observations K required to ensure recovery error $\leq \delta_L$ with a given reliability $1 - \epsilon$, $\epsilon \ll 1$, is within logarithmic in N, J, L factor off the “ideal” number of observations needed to achieve, with reliability $1 - \epsilon$, recovery error $\delta_L/2$.*

Recovering Linear Functionals on Unions of Convex Sets Revisited

Problem of interest: *Given*

- simple o.s. $\mathcal{O} = ((\Omega, \Pi), \{p_\mu(\cdot) : \mu \in \mathcal{M}\}, \mathcal{F})$
- collection $X_j, j = 1, \dots, J$, of nonempty convex compact sets in $\mathbb{R}^n$
- affine *sensing mappings* $A_j(\cdot)$ such that

$$A_j(x) \in \mathcal{M} \quad \forall x \in X_j$$

- linear function $g^T x$ on $\mathbb{R}^n$,
- stationary K -repeated observation $\omega^K = (\omega_1, \dots, \omega_K), \omega_k \sim p_{A_{j_x}(x)}$,
stemming from *unknown* signal $x \in X = \bigcup_j X_j$, such that $x \in X_{j_x}$

we want to estimate $g^T x$.

Note: When all sensing mappings $A_j(\cdot)$ are identical to each other, our problem becomes a special case of recovering 1-convex function on the union of convex sets. If, furthermore, $J = 1$ (there just one set $X = X_1$), we come back to the classical problem of linear functional estimation.

Risk: We quantify the performance of an estimate $\hat{g}(\omega^K)$ by its ϵ -risk:

$$\text{Risk}_\epsilon[\hat{g}|X] = \min \left\{ \rho : \text{Prob}_{\omega^K \sim p_{A_j(x)}^K} \left\{ |\hat{g}(\omega^K) - g^T x| > \rho \right\} \leq \epsilon \forall (j \leq J, x \in X_j) \right\}$$

— the smallest guaranteed size of $(1 - \epsilon)$ -confidence interval of the estimate.

- $\epsilon \in (0, 1)$: once for ever fixed tolerance.

We consider estimates which are aggregates of “pairwise” estimates of the form

$$\phi_{ij}^{(K)}(\omega^K) = \sum_k \phi_{ij}(\omega_k)$$

with $\phi_{ij} \in \mathcal{F}$.

Question: Given $\phi \in \mathcal{F}$, $\delta \in (0, 1)$, $\rho > 0$, and i, j , how to certify that

$$\text{Prob}_{\omega^K \sim p_{A_i(x)}^K} \left\{ \sum_k \phi(\omega_k) > g^T x + \rho \right\} \leq \delta \quad \forall x \in X_i \quad (a)$$

$$\text{Prob}_{\omega^K \sim p_{A_j(y)}^K} \left\{ \sum_k \phi(\omega_k) < g^T y - \rho \right\} \leq \delta \quad \forall y \in X_j \quad (b)$$

Answer: Define the “aggregate”

$$\Phi_{\mathcal{O}}(\phi, \mu) = \ln \left(\int \exp\{\phi(\omega)\} p_{\mu}(\omega) \Pi(d\omega) \right) : \mathcal{F} \times \mathcal{M} \rightarrow \mathbb{R}$$

By construction, $\Phi_{\mathcal{O}}(\cdot, \cdot)$ is convex in ϕ and concave in μ .

A sufficient condition for (a) and (b) is *existence of $\alpha > 0$ such that*

$$\max_{x \in X_i} \left[K\alpha \Phi_{\mathcal{O}}(\phi/\alpha, A_i(x)) - g^T x + \alpha \ln(1/\delta) \right] \leq \rho \quad (c)$$

$$\max_{y \in X_j} \left[K\alpha \Phi_{\mathcal{O}}(-\phi/\alpha, A_j(y)) + g^T y + \alpha \ln(1/\delta) \right] \leq \rho \quad (d)$$

Indeed when (c) takes place and $x \in X_i$, we have

$$\begin{aligned}
\pi &:= \text{Prob}_{\omega^K \sim p_{A_i(x)}^K} \left\{ \sum_k \phi(\omega_k) > g^T x + \rho \right\} \\
&\leq \mathbf{E}_{\omega^K \sim p_{A_i(x)}^K} \left\{ \exp \left(\frac{1}{\alpha} \left[\sum_k \phi(\omega_k) - g^T x - \rho \right] \right) \right\} \\
&= \exp \left(-[g^T x + \rho]/\alpha \right) \left[\mathbf{E}_{\omega \sim p_{A_i(x)}} \{ \exp(\phi(\omega)/\alpha) \} \right]^K \\
&= \exp([g^T x + \rho]/\alpha) \exp\{K \Phi_{\mathcal{O}}(\phi/\alpha, A_i(x))\} \\
&\Rightarrow \ln(\pi) \leq K \Phi_{\mathcal{O}}(\phi/\alpha, A_i(x)) - \alpha^{-1}[g^T x + \rho] \\
&\Rightarrow \alpha \ln(\pi) \leq K \alpha \Phi_{\mathcal{O}}(\phi/\alpha, A_i(x)) - g^T x - \rho \\
&\Rightarrow [\text{by (c)}] \alpha \ln(\pi) \leq -\alpha \ln(1/\delta) \Rightarrow \pi \leq \delta
\end{aligned}$$

Thus, (c) implies (a); in the same way, (d) implies (b).

$$\mathcal{O} = ((\Omega, \Pi), \{p_\mu(\cdot) : \mu \in \mathcal{M}\}, \mathcal{F})$$

$$\Phi_{\mathcal{O}}(\phi, \mu) = \ln \left(\int \exp\{\phi(\omega)\} p_\mu(\omega) \Pi(d\omega) \right) : \mathcal{F} \times \mathcal{M} \rightarrow \mathbb{R}$$

$$\begin{cases} \max_{x \in X_i} [K\alpha \Phi_{\mathcal{O}}(\phi/\alpha, A_i(x)) - g^T x + \alpha \ln(1/\delta)] \leq \rho & (c) \\ \max_{y \in X_j} [K\alpha \Phi_{\mathcal{O}}(-\phi/\alpha, A_j(y)) + g^T y + \alpha \ln(1/\delta)] \leq \rho & (d) \end{cases}$$

$$\Rightarrow \begin{cases} \text{Prob}_{\omega^K \sim p_{A_i(x)}^K} \left\{ \sum_k \phi(\omega_k) > g^T x + \rho \right\} \leq \epsilon/2 \quad \forall x \in X_i & (a) \\ \text{Prob}_{\omega^K \sim p_{A_j(y)}^K} \left\{ \sum_k \phi(\omega_k) < g^T y - \rho \right\} \leq \epsilon/2 \quad \forall y \in X_j & (b) \end{cases}$$

Observation: Let $\rho > 0$, $\psi \in \mathcal{F}$ and $\alpha > 0$ be such that

$$\max_{x \in X_i, y \in X_j} \frac{1}{2} [K\alpha [\Phi_{\mathcal{O}}(\psi/\alpha, A_i(x)) + \Phi_{\mathcal{O}}(-\psi/\alpha, A_j(y))] + g^T [y - x]] + \alpha \ln(1/\delta) \leq \rho \quad (*)$$

Then, setting

$$\varkappa = \frac{1}{2K} \left[\underbrace{\max_{y \in X_j} [K\alpha \Phi_{\mathcal{O}}(-\psi/\alpha, A_j(y)) + g^T y]}_{\gamma_-} - \underbrace{\max_{x \in X_i} [K\alpha \Phi_{\mathcal{O}}(\psi/\alpha, A_i(x)) - g^T x]}_{\gamma_+} \right],$$

$$\phi(\cdot) = \psi(\cdot) + \varkappa$$

so that $\phi \in \mathcal{F}$, we ensure (c) and (d) and thus ensure (a), (b).

Indeed, (*) reads $\frac{1}{2}[\gamma_+ + \gamma_-] + \alpha \ln(1/\delta) \leq \rho$, and clearly

$$\Phi_{\mathcal{O}}(\phi/\alpha, \cdot) = \Phi_{\mathcal{O}}(\psi/\alpha, \cdot) + \varkappa/\alpha, \quad \Phi_{\mathcal{O}}(-\phi/\alpha, \cdot) = \Phi_{\mathcal{O}}(-\psi/\alpha, \cdot) - \varkappa/\alpha,$$

whence

$$\begin{aligned} \max_{x \in X_i} & \left[K\alpha \Phi_{\mathcal{O}}(\phi/\alpha, A_i(x)) - g^T x + \alpha \ln(1/\delta) \right] \\ &= \gamma_+ + K\varkappa + \alpha \ln(1/\delta) = \gamma_+ + \frac{1}{2}[\gamma_- - \gamma_+] + \alpha \ln(1/\delta) \\ &= \frac{1}{2}[\gamma_+ + \gamma_-] + \alpha \ln(1/\delta) \leq \rho, \end{aligned}$$

and

$$\begin{aligned} \max_{y \in X_j} & \left[K\alpha \Phi_{\mathcal{O}}(-\phi/\alpha, A_i(x)) + g^T y + \alpha \ln(1/\delta) \right] \\ &= \gamma_- - K\varkappa + \alpha \ln(1/\delta) = \gamma_- - \frac{1}{2}[\gamma_- - \gamma_+] + \alpha \ln(1/\delta) \\ &= \frac{1}{2}[\gamma_+ + \gamma_-] + \alpha \ln(1/\delta) \leq \rho \end{aligned}$$

Building the estimate: Consider convex optimization problems

$$\text{Opt}_{ij}(K) = \min_{\alpha > 0, \psi \in \mathcal{F}} \left\{ \max_{x \in X_i, y \in X_j} \frac{1}{2} [K\alpha [\Phi_{\mathcal{O}}(\psi/\alpha, A_i(x)) + \Phi_{\mathcal{O}}(-\psi/\alpha, A_i(y))] + g^T[y - x]] + \alpha \ln(2J/\epsilon) \right\}.$$

Let α_{ij}, ψ_{ij} be feasible solutions to the problems, and ρ_{ij} be the corresponding values of the objectives (the less, the better).

• We set

$$\begin{aligned} \varkappa_{ij} &= \frac{1}{2K} \left[\max_{y \in X_j} [K\alpha_{ij}\Phi_{\mathcal{O}}(-\psi_{ij}/\alpha_{ij}, A_j(y)) + g^T y] \right. \\ &\quad \left. - \max_{x \in X_i} [K\alpha_{ij}\Phi_{\mathcal{O}}(\psi_{ij}/\alpha_{ij}, A_i(x)) - g^T x] \right], \\ \phi_{ij}(\cdot) &= \psi_{ij}(\cdot) + \varkappa_{ij}, \quad \phi_{ij}^{(K)}(\omega^K) = \sum_k \phi_{ij}(\omega_k), \end{aligned}$$

thus ensuring

$$\begin{aligned} \text{Prob}_{\omega^K \sim p_{A_i(x)}^K} \left\{ \phi_{ij}^{(K)}(\omega^K) > g^T x + \rho_{ij} \right\} &\leq \frac{\epsilon}{2J}, \quad x \in X_i \\ \text{Prob}_{\omega^K \sim p_{A_j(y)}^K} \left\{ \phi_{ij}^{(K)}(\omega^K) < g^T y - \rho_{ij} \right\} &\leq \frac{\epsilon}{2J}, \quad y \in X_j \end{aligned}$$

$$\text{Prob}_{\omega^K \sim p_{A_i(x)}^K} \left\{ \phi_{ij}^{(K)}(\omega^K) > g^T x + \rho_{ij} \right\} \leq \frac{\epsilon}{2J}, x \in X_i$$

$$\text{Prob}_{\omega^K \sim p_{A_j(y)}^K} \left\{ \phi_{ij}^{(K)}(\omega^K) < g^T y - \rho_{ij} \right\} \leq \frac{\epsilon}{2J}, y \in X_j$$

- Given observation ω^K , we build the matrix

$$\left[E_{ij} := \phi_{ij}^{(K)}(\omega^K) \right]_{i,j \leq J},$$

set

$$r_i = \max_j E_{ij}, \quad r = \min_i r_i = \min_i \max_j E_{ij},$$

$$c_j = \min_i E_{ij}, \quad c = \max_j c_j = \max_j \min_i E_{ij},$$

and specify the estimate $\hat{g}(\omega^K)$ of $g^T x$, $x \in X = \cup_j X_j$, as

$$\hat{g}(\omega^K) = \frac{1}{2}[r + c].$$

Theorem ϵ -Risk of the estimate $\hat{g}$ satisfies

$$\text{Risk}_\epsilon[\hat{g}|X] \leq \rho := \max_{i,j} \rho_{ij}.$$

$$\text{Prob}_{\omega^K \sim p_{A_i(x)}^K} \left\{ \phi_{ij}^{(K)}(\omega^K) > g^T x + \rho_{ij} \right\} \leq \frac{\epsilon}{2J}, x \in X_i \quad (a)$$

$$\text{Prob}_{\omega^K \sim p_{A_j(y)}^K} \left\{ \phi_{ij}^{(K)}(\omega^K) < g^T y - \rho_{ij} \right\} \leq \frac{\epsilon}{2J}, y \in X_j$$

$$E_{ij} = \phi_{ij}^{(K)}(\omega^K) \quad (b)$$

$$r_i = \max_j E_{ij}, \quad r = \min_i r_i = \min_i \max_j E_{ij}, \quad c_j = \min_i E_{ij}, \quad c = \max_j c_j = \max_j \min_i E_{ij} \quad (c)$$

$$\widehat{g}(\omega^K) = \frac{1}{2}[r + c] \quad (d)$$

Proof. 1⁰. By (b),

$$r_i \geq E_{ij} \geq c_j \Rightarrow r_i \geq c_j \quad \forall i, j \quad (e)$$

2⁰. Let observation ω^K stem from signal $x \in X_\ell$, let P be the distribution of ω^K , and let $\mathcal{E}$ be the event

$$\phi_{\ell j}^{(K)}(\omega^K) \leq g^T x + \rho_{\ell j} \quad \forall j \leq J \quad \text{and} \quad \phi_{i\ell}^{(K)}(\omega^K) \geq g^T x - \rho \quad \forall i \leq J,$$

so that by (a) and union bound,

$$P(\mathcal{E}) \geq 1 - \epsilon.$$

Let $\Delta = [g^T x - \rho, g^T x + \rho]$. All we need to verify is that *when $\omega^K \in \mathcal{E}$, one has $r, c \in \Delta$.*

Indeed, when $\omega^K \in \mathcal{E}$, we have

$$r_{\ell j} \leq g^T x + \rho_{\ell j} \quad \forall j \Rightarrow r_\ell \leq g^T x + \rho \Rightarrow r \leq g^T x + \rho,$$

same as

$$c_{i\ell} \geq g^T x - \rho_{i\ell} \quad \forall i \Rightarrow c_\ell \geq g^T x - \rho \Rightarrow c \geq g^T x - \rho.$$

Since $r \geq c$ by (e), we conclude that $r, c \in \Delta$ whenever $\omega^K \in \mathcal{E}$. □

Near-Optimality

$$\text{Opt}_{ij}(K) = \min_{\alpha > 0, \psi \in \mathcal{F}} \left\{ \max_{x \in X_i, y \in X_j} \frac{1}{2} [K\alpha [\Phi_{\mathcal{O}}(\psi/\alpha, A_i(x)) + \Phi_{\mathcal{O}}(-\psi/\alpha, A_i(y))] + g^T[y - x]] + \alpha \ln(2J/\epsilon) \right\}. \quad (P_{ij})$$

Note: Every $\text{Opt}_{ij}(K)$ is either a real, or $-\infty$. It is immediately seen that $\text{Opt}_{ii}(K)$ is a nonnegative real, so that the quantity

$$\text{Opt}(K) = \max_{i,j} \text{Opt}_{ij}(K)$$

is a nonnegative real. Properly selecting feasible solutions to the convex optimization problems (P_{ij}) , we can make the upper bound ρ on ϵ -risk of our estimate arbitrarily close to $\text{Opt}(K)$.

We are about to demonstrate that $\text{Opt}(K)$ *is nearly the optimal ϵ -risk* in our estimation problem

$$\text{RiskOpt}_{\epsilon}(K) = \inf_{\hat{g}} \text{Risk}_{\epsilon}[\hat{g}|X],$$

the infimum being taken over *all* K -observation estimates.

Theorem Let $\epsilon \in (0, 1/2)$ and M be a positive integer. Then

(i) for every integer K satisfying

$$\frac{K}{M} > \frac{2 \ln(2J/\epsilon)}{\ln([4\epsilon(1 - \epsilon)]^{-1})} \quad [\rightarrow 2, \text{ as } \epsilon \rightarrow +0]$$

one has

$$\text{Opt}(K) \leq \text{RiskOpt}_\epsilon(M).$$

(ii) In addition, in the special case where for every i, j there exists $\bar{x}_{ij} \in X_i \cap X_j$ such that $A_i(\bar{x}_{ij}) = A_j(\bar{x}_{ij})$ one has

$$K \geq M \Rightarrow \text{Opt}(K) \leq \frac{2 \ln(2J/\epsilon)}{\ln([4\epsilon(1 - \epsilon)]^{-1})} \text{RiskOpt}_\epsilon(M).$$

$$\text{Opt}_{ij}(K) = \inf_{\alpha > 0, \psi \in \mathcal{F}} \left\{ \max_{x \in X_i, y \in X_j} \frac{1}{2} [K\alpha [\Phi_{\mathcal{O}}(\psi/\alpha, A_i(x)) + \Phi_{\mathcal{O}}(-\psi/\alpha, A_i(y))] + g^T[y - x]] + \alpha \ln(2J/\epsilon) \right\}$$

Proof of (i): Let us fix i and j . $\text{Opt}_{ij}(K)$ is the saddle point value in the saddle point problem

$$\inf_{\alpha > 0, \psi \in \mathcal{F}} \max_{(x, y) \in X_i \times X_j} \left[\frac{1}{2} K\alpha [\Phi_{\mathcal{O}}(\psi/\alpha, A_i(x)) + \Phi_{\mathcal{O}}(-\psi/\alpha, A_i(y))] + \frac{1}{2} g^T[y - x] + \alpha \ln(2J/\epsilon) \right]$$

which is convex-concave with compact maximization domain
 $\Rightarrow$ [by Sion-Kakutani Theorem]

$$\text{Opt}_{ij}(K) = \max_{(x, y) \in X_i \times X_j} Q_{ijK}(x, y),$$

$$Q_{ijK}(x, y) = \inf_{\alpha > 0, \psi \in \mathcal{F}} \left[\frac{1}{2} K\alpha [\Phi_{\mathcal{O}}(\psi/\alpha, A_i(x)) + \Phi_{\mathcal{O}}(-\psi/\alpha, A_i(y))] + \alpha \ln(2J/\epsilon) \right] + \frac{1}{2} g^T[y - x]$$

Setting $\phi = \psi/\alpha$, $\mu = A_i(x)$, $\nu = A_j(y)$, and recalling what $\Phi_{\mathcal{O}}$ is, we get

$$Q_{ijK}(x, y) = \inf_{\alpha > 0} \alpha \left[\ln(2J/\epsilon) + \inf_{\phi \in \mathcal{F}} \frac{K}{2} \left[\ln \left(\int e^{\phi(\omega)} p_{\mu}(\omega) \Pi(d\omega) \right) + \ln \left(\int e^{-\phi(\omega)} p_{\nu}(\omega) \Pi(d\omega) \right) \right] \right] + \frac{1}{2} g^T[y - x].$$

$$Q_{ijK}(x, y) = \inf_{\alpha > 0} \alpha \left[\ln(2J/\epsilon) + \inf_{\phi \in \mathcal{F}} \frac{K}{2} \left[\ln \left(\int e^{\phi(\omega)} p_\mu(\omega) \Pi(d\omega) \right) + \ln \left(\int e^{-\phi(\omega)} p_\nu(\omega) \Pi(d\omega) \right) \right] \right] + \frac{1}{2} g^T [y - x]$$

We claim that *the inner* $\inf_{\phi}$ *is attained at*

$$\phi(\omega) = \phi_*(\omega) := \frac{1}{2} \ln(p_\nu(\omega)/p_\mu(\omega))$$

(note that $\phi_* \in \mathcal{F}$ since $\mathcal{O}$ is simple).

Indeed, we have for $\phi(\omega) = \phi_*(\omega) + \delta(\omega)$:

$$\begin{aligned} & \inf_{\phi \in \mathcal{F}} \frac{K}{2} \left[\ln \left(\int \exp\{\phi(\omega)\} p_\mu(\omega) \Pi(d\omega) \right) + \ln \left(\int \exp\{-\phi(\omega)\} p_\nu(\omega) \Pi(d\omega) \right) \right] \\ &= \inf_{\delta \in \mathcal{F}} \frac{K}{2} \left[\ln \left(\int \exp\{\phi_*(\omega) + \delta(\omega)\} p_\mu(\omega) \Pi(d\omega) \right) + \ln \left(\int \exp\{-\phi_*(\omega) - \delta(\omega)\} p_\nu(\omega) \Pi(d\omega) \right) \right] \\ &= \inf_{\delta \in \mathcal{F}} \underbrace{\left[\ln \left(\int \exp\{\delta(\omega)\} \sqrt{p_\mu(\omega)p_\nu(\omega)} \Pi(d\omega) \right) + \ln \left(\int \exp\{-\delta(\omega)\} \sqrt{p_\mu(\omega)p_\nu(\omega)} \Pi(d\omega) \right) \right]}_{f(\delta)} \end{aligned}$$

$f(\delta)$ is convex and even function of $\delta \in \mathcal{F}$ and as such achieves its minimum at $\delta = 0$.

- By the above computation,

$$\begin{aligned}
\text{Opt}_{ij}(K) &= \max_{x \in X_i, y \in X_j} Q_{ijK}(x, y) \\
&= \max_{x \in X_i, y \in X_j} \left[\frac{1}{2} g^T[y - x] + \inf_{\alpha > 0} \alpha \left[K \ln \left(\int \sqrt{p_{A_i(x)}(\omega) p_{A_j(y)}(\omega)} \Pi(d\omega) \right) + \ln(2J/\epsilon) \right] \right] \\
&= \max_{x, y} \left\{ \frac{1}{2} g^T[y - x] : K \ln \left(\int \sqrt{p_{A_i(x)}(\omega) p_{A_j(y)}(\omega)} \Pi(d\omega) \right) + \ln(2J/\epsilon) \geq 0, x \in X_i, y \in X_j \right\}
\end{aligned}$$

- To complete the proof, it suffices to verify that

$$\frac{K}{M} > \frac{2 \ln(2J/\epsilon)}{\ln([4\epsilon(1 - \epsilon)]^{-1})} \Rightarrow \text{Opt}_{ij}(K) \leq \text{RiskOpt}_{\epsilon}(M) \text{ for all } i, j.$$

Indeed, assuming the opposite, we can find i, j and $\bar{x} \in X_i, \bar{y} \in X_j$ such that setting $p = p_{A_i(\bar{x})}, q = p_{A_j(\bar{y})}$, we have

$$\begin{aligned}
\frac{1}{2} g^T[\bar{y} - \bar{x}] &> \text{RiskOpt}_{\epsilon}(M) & (a) \\
K \ln \left(\int \sqrt{p(\omega) q(\omega)} \Pi(d\omega) \right) + \ln(2J/\epsilon) &\geq 0 & (b)
\end{aligned}$$

$$\frac{1}{2}g^T[\bar{y} - \bar{x}] > \text{RiskOpt}_\epsilon(M) \quad (a)$$

$$K \ln \left(\int \sqrt{p(\omega)q(\omega)} \Pi(d\omega) \right) + \ln(2J/\epsilon) \geq 0 \text{ for some } \bar{x} \in X_i, \bar{y} \in X_j \quad (b)$$

- (a) implies that simple hypotheses

$$H_1 : \omega^M \sim p^M, \quad H_2 : \omega^M \sim q^M$$

about distribution of ω^M can be decided upon with risk $\leq \epsilon$, whence

$$\int \min[p^M(\omega^M), q^M(\omega^M)] \Pi^M(d\omega^M) \leq 2\epsilon$$

$$\begin{aligned} \Rightarrow \left[\int \sqrt{p(\omega)q(\omega)} \Pi(d\omega) \right]^M &= \int \sqrt{p^M(\omega^M), q^M(\omega^M)} \Pi^M(d\omega^M) \\ &= \int \sqrt{\min[p^M(\omega^M), q^M(\omega^M)]} \sqrt{\max[p^M(\omega^M), q^M(\omega^M)]} \Pi^M(d\omega^M) \\ &\leq \left[\int \min[p^M(\omega^M), q^M(\omega^M)] \Pi^M(d\omega^M) \right]^{1/2} \left[\int \max[p^M(\omega^M), q^M(\omega^M)] \Pi^M(d\omega^M) \right]^{1/2} \\ &= \left[\int \min[p^M(\omega^M), q^M(\omega^M)] \Pi^M(d\omega^M) \right]^{1/2} \left[\int \max[p^M(\omega^M), q^M(\omega^M)] \Pi^M(d\omega^M) \right]^{1/2} \\ &= \left[\int \min[p^M(\omega^M), q^M(\omega^M)] \Pi^M(d\omega^M) \right]^{1/2} \\ &\quad \times \left[\int [p^M(\omega^M) + q^M(\omega^M) - \min[p^M(\omega^M), q^M(\omega^M)]] \Pi^M(d\omega^M) \right]^{1/2} \\ &\leq \sqrt{2\epsilon(2 - 2\epsilon)} \\ \Rightarrow \left[\int \sqrt{p(\omega)q(\omega)} \Pi(d\omega) \right]^K &\leq \left[2\sqrt{\epsilon(1 - \epsilon)} \right]^{K/M} < \frac{\epsilon}{2J} \end{aligned}$$

[due to $\frac{K}{M} > \frac{2 \ln(2J/\epsilon)}{\ln([4\epsilon(1-\epsilon)]^{-1})}$ and $\epsilon < 1/2$], what contradicts (b). □

Illustration: Estimating Cumulative Distribution from Noisy Observations

- Let X being convex and closed subset of Δ_n^o , and let $\zeta \sim x \in X$ be a discrete random variable on $\{1, \dots, n\}$,
(i.e. $P(\zeta = i) = x_i, i = 1, \dots, n$).
- Let $\omega \in \{1, \dots, d\}$ be a discrete random variable, $\omega \sim \mu = Ax$, where $A \in \mathbb{R}_{++}^{d \times n}$ is a given stochastic matrix.

Typical examples are

- Convolution model $\omega = \zeta + \xi$ where ξ (the noise) is a discrete random variable with known distribution
- Multiplicative noise model $\omega = \xi \zeta$
- Missing observations model
- ...

Our objective is, given a stationary K -repeated observation $\omega^K = [\omega_1 \dots \omega_K]$, estimate values $F_i(x) = e_i^T x := \sum_{j=1}^i x_j$ of the c.d.f. of x .

Note: We are in the Discrete o.s. and can apply our apparatus to estimate the value of linear functionals $F_i(x)$

To compute the estimator $\hat{F}_i(\cdot)$ of F_i we proceed as follows.

- We find an optimal solution $[\bar{x}, \bar{y}]$ to the problem

$$\max \left\{ \frac{1}{2} e^T (y - x) : x, y \in X, K \ln \left[\sum_{i=1}^d \sqrt{[Ax]_i [Ay]_i} \right] + \ln(2/\epsilon) \geq 0 \right\},$$

same as solving an equivalent *conic quadratic problem*

$$\max_{x, y \in X} \left\{ \frac{1}{2} e^T (y - x) : x, y \in X, \sum_{i=1}^d \sqrt{[Ax]_i [Ay]_i} \geq (2/\epsilon)^{1/K} \right\}. \quad (*)$$

- Let $[\bar{x}; \bar{y}]$ be an optimal solution to $(*)$. We set

$$\hat{F}_i(\omega^K) = \phi_*^T \hat{p}^K + c_*, \quad [\phi_*]_i = \nu \left(\sum_{i=1}^M \sqrt{[A\bar{x}]_i [A\bar{y}]_i} \right) \ln \sqrt{\frac{[A\bar{y}]_i}{[A\bar{x}]_i}}, \quad c_* = \frac{1}{2} g^T (\bar{x} + \bar{y}),$$

where $\hat{p}^K$ is the empirical distribution of ω :

$$\hat{p}_i^K = \frac{1}{K} \sum_{k=1}^K 1_{\omega_k=i},$$

and ν is the Lagrange multiplier of the conic constraint.

Theorem Let $\epsilon \in (0, 1/2]$. The risk of the estimator $\hat{g}_*$ satisfies:

$$\text{Risk}_\epsilon[\hat{F}_i|X] \leq \frac{2 \ln(2/\epsilon)}{\ln([4\epsilon(1-\epsilon)]^{-1})} \text{RiskOpt}_\epsilon(K). \quad [\rightarrow 2, \text{ as } \epsilon \rightarrow +0]$$

Application: Convolution Model

Goal: Suppose that r.v. $\eta \sim F_\eta$, $\text{supp } \eta \subseteq [-3, 3]$, $v \sim F_v$, with known F_v . Given $n = 100$ i.i.d. observations Y_i of a r.v. $\eta + v$, we want to recover $F_\eta(t)$ for a given $t \in [-3, 3]$.

- We split $[-3, 3]$ into $n = 60$ equal segments δ_i of length δ , with the “sampled signal” $x_i = P(\eta \in \Delta_i)$. On suppose

$$\frac{x_i - 2x_{i-1} + x_{i-2}}{\Delta^2} \leq 1$$

(a “discrete approximation” of *the density f_η of η is twice differentiable with $|f''_\eta(t)| \leq 1$*).

- We cut the domain of Y into $d = 42$ chunks γ_i (40 equal segments between $]Y_{\min}, Y_{\max}[$ and two “complements” in $\mathbb{R}$). The distribution of $\eta + v$ is approximated with $\mu = Ax$, where

$$A_{ij} = \delta_i^{-1} \int_{v \in \gamma_i} \int_{u \in \delta_j} f_\xi(v - u) dv du.$$

Numerical illustration: Estimation of $F_\eta(0)$, normal nuisance v , $n = 100$ observations, $e = 0.05$.

Optimal signals $\bar{x}$ et $\bar{y}$

Corresponding c.d.f. of $A\bar{x}$ and of $A\bar{y}$

— risk of minmax affine estimator
— constrained to $E\{\eta\} = 0$ risk
— symmetry constrained risk

— Gaussian noise v
— Laplace noise v

Sequential Hypothesis Testing

Motivating example: Election Polls. Population-wide elections with L candidates are to be held.

Preferences are represented by L -dimensional basic orth

$$[0; \dots; 0; 1; 0; \dots; 0]$$

with 1 in position ℓ meaning *the vote is for candidate ℓ* .

Equivalently: Preference ω of a voter is a vertex in the probabilistic simplex

$$\Delta_L = \{p \in \mathbb{R}^L : p \geq 0, \sum_{\ell} p_{\ell} = 1\}.$$

- The average $\mu = [\mu_1; \dots; \mu_L]$ of preferences of all voters “encodes” election’s outcome: μ_{ℓ} is the fraction of voters supporting ℓ -th candidate, and the winner corresponds to the largest entry in μ (assumed to be uniquely defined).

μ is a probabilistic vector: $\mu \in \Delta_L$. We think of μ as of a probability distribution on the L -element set $\Omega = \text{Ext}(\Delta_L)$ of vertices of Δ_L .

Our goal is to design *election poll* – to select K voters at random from the uniform over the population distribution and to observe their preferences, in order to predict, with reliability $1 - \epsilon$, election's outcome.

The model is drawing a stationary K -repeated observation $\omega^K = (\omega_1, \dots, \omega_K)$, $\omega_k \in \Omega$, from distribution μ .

Assume *once for ever that the elections never end with “near tie,” that is, the fraction of votes for the winner is at least by a known margin δ larger than the fraction of votes for every no-winner.*

We introduce L hypotheses on the distribution μ of $\omega_1, \dots, \omega_K$:

$$H_\ell : \mu \in \mathcal{P}_\ell = \{\mu \in \Delta_L : \mu_\ell \geq \mu_{\ell'} + \delta, \forall \ell' \neq \ell\}, \quad \ell = 1, \dots, L$$

- Our goal is to specify K in a way which allows to decide on $H_1, \dots, H_L$ via stationary K -repeated observations *with risk $\leq \epsilon$.*
- We are in the case of Discrete o.s., and can use our machinery to build a near-optimal K -observation test deciding on $H_1, \dots, H_L$ up to trivial closeness

$$\mathcal{C} : \text{“} H_\ell \text{ is close to } H_{\ell'} \text{ iff } \ell = \ell' \text{”}$$

and select the smallest K for which the $\mathcal{C}$ -risk of this test is $\leq \epsilon$.

Illustration $L = 2$: In this case Ω is a two-point set of basic orths in $\mathbb{R}^2$, the minimum risk single-observation detector is

$$\phi_*(\omega) = \frac{1}{2} \ln \left(\frac{1 + \delta}{1 - \delta} \right) [\omega^1 - \omega^2] : \Omega \rightarrow \mathbb{R}$$

and $\text{Risk}[\phi_* | \mathcal{P}_1, \mathcal{P}_2] = 1 - \delta^2$

$$\Rightarrow K = \text{Ceil} \left(\frac{\ln(1/\epsilon)}{\ln(1/(1-\delta^2))} \right) \asymp \frac{1}{\delta^2} \ln(1/\epsilon).$$

δ	0.3162	0.1000	0.0316	0.0100
$\underline{K}, K$	51, 88	534, 917	5379, 9206	53820, 92064

$\underline{K}$: *lower* bound on optimal poll size, $\epsilon = .01$

USA Presidential Elections-2016:

State	Actual margin	Poll size, lower bound	Poll size, upper bound
Wisconsin	0.0041	24 576	88 663
Pennsylvania	0.0038	28 978	104 545
Michigan	0.0012	281 958	1 017 227

Confidence level 95%

$$K \asymp \frac{1}{\delta^2} \ln(1/\epsilon)$$

δ	0.3162	0.1000	0.0316	0.0100
$\underline{K}, \underline{K}$	51, 88	534, 917	5379, 9206	53820, 92064

$\underline{K}$: *lower* bound on optimal poll size, $\epsilon = .01$

Required size of election poll grows rapidly as “winning margin” decreases.

Question: *Can we do better?*

Partial remedy: Let us pass to *sequential tests*, where we *attempt* to make conclusion *before* all K respondents required by the worst-case-oriented analysis are interviewed.

If elections are about to be “landslide” (i.e., in unknown to us actual distribution μ_* of voters’ preferences the winner beats all other candidates by margin $\delta_* \gg \delta$), the winner hopefully can be identified after a relatively small number of interviews.

Strategy. We select a number S of *valuations* and associate with valuation s number $K(s)$ of observations, $K(1) < \dots < K(S)$.

- the s -th valuation is complete when $K(s)$ observations are collected; then, we apply to the collected so far observation $\omega^{K(s)} = (\omega_1, \dots, \omega_{K(s)})$ a test $\mathcal{T}_s$ which, depending on $\omega^{K(s)}$,
 - either accepts exactly one of the hypotheses $H_1, \dots, H_L$, in which case we terminate,
 - or claims that information collected so far is not sufficient to decide, in which case we either collect more observations (when $s < S$) or terminate (when $s = S$).

Specifications: We want the overall procedure to be

- *conclusive*: a decision should be made during one of the S valuations (thus, inference *should be guaranteed* when valuation S is reached);
- *reliable*: whenever the true distribution μ_* underlying observations obeys one of our L hypotheses, the μ_* -probability for this hypothesis to be accepted should be $\geq 1 - \epsilon$, where $\epsilon \in (0, 1)$ is a given in advance risk bound.

Implementation: We select somehow the number of valuations S and set $\delta_s = \delta^{s/S}$ so that $\delta_1 > \delta_2 > \dots > \delta_S = \delta$; we split risk bound ϵ into S parts ϵ_s : $\epsilon_s > 0$, $s \leq S$ & $\sum_{s=1}^S \epsilon_s = \epsilon$;

- For $s < S$, we define $2L$ hypotheses

$$H_{2\ell-1}^s = H_\ell : \mu \in \mathcal{P}_{2\ell-1}^s = \mathcal{P}_\ell := \{\mu \in \Delta_L : \mu_\ell \geq \delta_s + \max_{\ell' \neq \ell} \mu_{\ell'}\}$$

“weak hypothesis”

$$H_{2\ell}^s = \{\mu \in \mathcal{P}_{2\ell}^s := \{\mu \in \Delta_L : \mu_\ell \geq \delta_s + \max_{\ell' \neq \ell} \mu_{\ell'}\} \subset \mathcal{P}_\ell\}$$

“strong hypothesis”

$1 \leq \ell \leq L$, and assign $H_{2\ell-1}^s$ and $H_{2\ell}^s$ with color ℓ , $1 \leq \ell \leq L$.

- For $s = S$ we consider L hypotheses $H_\ell^S = H_\ell$, $\ell = 1, \dots, L$, with H_ℓ^S assigned color ℓ .
- For $s < S$, we introduce *closeness relation* $\mathcal{C}_s$ on the collection of hypotheses $H_1^s, \dots, H_{2L}^s$ as follows:
 - the only hypotheses close to a strong hypothesis $H_{2\ell}^s$ are the hypotheses $H_{2\ell}^s$ and $H_{2\ell-1}^s$ of the same color;
 - the only hypotheses close to a weak hypothesis $H_{2\ell-1}^s$ are all weak hypotheses and the strong hypothesis $H_{2\ell}^s$ of the same color as $H_{2\ell-1}^s$.
 - For $s = S$, the $\mathcal{C}_s$ -closeness is trivial: $H_\ell^S \equiv H_\ell$ is $\mathcal{C}_S$ -close to $H_{\ell'}^S \equiv H_{\ell'}$ if and only if $\ell = \ell'$.

3-candidate hypotheses in the probabilistic simplex Δ_3

[weak green]	M_1	dark green + light green: candidate A wins with margin $\geq \delta_S$
[strong green]	M_1^s	dark green: candidate A wins with margin $\geq \delta_s > \delta_S$
[weak red]	M_2	dark red + pink: candidate B wins with margin $\geq \delta_S$
[strong red]	M_2^s	dark red: candidate B wins with margin $\geq \delta_s > \delta_S$
[weak blue]	M_3	dark blue + light blue: candidate C wins with margin $\geq \delta_S$
[strong blue]	M_3^s	dark blue: candidate C wins with margin $\geq \delta_s > \delta_S$

- $H_{2\ell-1}^s : \mu \in M_\ell$ [weak hypothesis]
weak hypothesis $H_{2\ell-1}^s$ is $\mathcal{C}_s$ -close to itself, to all other weak hypotheses and to strong hypothesis $H_{2\ell}^s$ of the same color as $H_{2\ell-1}^s$
- $H_{2\ell}^s : \mu \in M_\ell^s$ [strong hypothesis]
strong hypothesis $H_{2\ell}^s$ is $\mathcal{S}$ -close only to itself and to weak hypothesis $H_{2\ell-1}^s$ of the same color as $H_{2\ell-1}^s$

Note: We are in the case of stationary repeated observations in Discrete o.s., and the hypotheses H_j^s are of the form

“ i.i.d. observations $\omega_1, \omega_2, \dots$ are drawn from distribution $\mu \in M_j^s$ with nonempty closed convex sets $M_j^s \subset \Delta_L$ ”, *and sets $M_j^s, M_{j'}^s$ with $(j, j') \notin \mathcal{C}_s$ do not intersect*

$\Rightarrow$ risks of the optimal pairwise detectors for $\mathcal{P}_j^s, \mathcal{P}_{j'}^s, (j, j') \notin \mathcal{C}_s$, are < 1

$\Rightarrow$ we can efficiently find the smallest $K = K(s)$ for which a test $\mathcal{T} = \mathcal{T}_s$ deciding, via stationary $K(s)$ -repeated observations, on the hypotheses $\{H_j^s\}_j$ has $\mathcal{C}_s$ -risk $\leq \epsilon_s$.

Our inference routine works as follows:

- we observe ω_k , $k = 1, 2, \dots, K(S)$ (i.e., carry interviews with one by one randomly selected voters), and perform s -th valuation to make conclusion when $K(s)$ observations are acquired ($K(s)$ interviews are completed).
- At s -th valuation, we apply the test $\mathcal{T}_s$ to observation $\omega^{K(s)}$. If the test does accept some of the hypotheses H_j^s *and all accepted hypotheses have the same color ℓ* , we accept ℓ -th of our original hypotheses $H_1, \dots, H_L$ (i.e., predict that ℓ -th candidate will be the winner) and terminate,
- otherwise we proceed to next observation round (i.e., next interviews) (when $s < S$) or claim the winner to be, say, the first candidate and terminate (when $s = S$).

- The risk of the outlined sequential hypothesis testing procedure is $\leq \epsilon$: whenever the distribution μ_* underlying observations obeys hypothesis H_ℓ for some $\ell \leq L$, the μ_* -probability of the event “ H_ℓ is the only accepted hypothesis” is at least $1 - \epsilon$.
- The worst-case volume of observations $K(S)$ is within logarithmic factor from the minimal number of observations allowing to decide on the hypotheses $H_1, \dots, H_L$ with risk $\leq \epsilon$.
- Whenever the distribution μ_* underlying observations obeys strong hypothesis $H_{2\ell}^s$ for some ℓ and s (“distribution μ_* of voters’ preferences corresponds to winning margin at least δ_s ”), the conclusion, with μ_* -probability $\geq 1 - \epsilon$, will be made in course of the first s valuations (i.e., in course of the first $K(s)$ interviews).

Informally: In “landslide elections,” the winner will be predicted reliably after a small number of interviews.

Numerical Illustration: 2-Candidate Elections

Setup:

- # of candidates $L = 2$
- $\delta_s = 10^{-s/4}$
- range of # of valuations S : $1 \leq S \leq 8$

Numerical Results:

S	1	2	4	5	6	8
$\delta = \delta_S$	0.5623	0.3162	0.1000	0.0562	0.0316	0.0100
K	25	88	287	917	9206	92098
$K(S)$	25	152	1594	5056	16005	160118

Volume K of non-sequential test, number of valuations S and worst-case volume $K(S)$ of sequential test as functions of winning margin $\delta = \delta_S$. Risk ϵ is set to 0.01.

Note that worst-case volume of sequential test is essentially worse than the volume of non-sequential test.

But: When drawing the true distribution μ_* of voters' preferences at random from the uniform distribution on the set of μ 's with winning margin ≥ 0.01 , the typical size of observations used by Sequential test with $S = 8$ prior to termination is $\ll K(S)$:

Empirical Volume of Sequential test

median	mean	60%	65%	75%	80%	85%	90%	95%	100%
177	9182	177	397	617	1223	1829	8766	87911	160118

Column "X%": empirical X%-quantile of test's volume. Data over 1,000 experiments.
Empirical risk: 0.01

Measurement Design

Observation: In our *Hypothesis Testing setup*, observation scheme is our “environment” and is completely out of our control. However, there are situations where the observation scheme is under our *partial* control.

Example: Election Poll revisited. In our original problem, a particular voter was represented by basic orth $\omega = [0; \dots; 0; 1; 0; \dots; 0] \in \mathbb{R}^L$, with entry 1 in position ℓ meaning that the voter prefers candidate ℓ to all other candidates. Our goal was to predict the winner by observing preferences of respondents selected at random from uniform distribution on voters’ population.

A modification: Imagine we can split voters in I non-intersecting groups (say, according to age, education, gender, income, occupation,...) in such a way that we have certain a priori knowledge of the distributions of preferences within the group.

In this situation, our poll can be reorganized:

- We assign groups with nonnegative weights q_i summing up to 1
- To organize an interview, we first select at random one of the groups, with probability q_i to select group i , and then select a respondent from i -th group at random, from uniform distribution on the group.

Note: When q_i is equal to the fraction θ_i of group i in the entire population, the above policy reduces to the initial one. It can make sense, however, to use q_i different from θ_i , with $q_i \ll \theta_i$ if a priori information about preferences of voters from i -th group is rich, and $q_i \gg \theta_i$ if this information is poor. Hopefully, this will allow us to make more reliable predictions with the same total number of interviews.

The model of outlined situation is as follows:

- We characterize distribution of preferences within group i by vector $\mu^i \in \Delta_L$. for $1 \leq j \leq L$, j -th entry in μ^i is the fraction of voters *in group i* voting for candidate j ;

The population-wide distribution of voters' preferences is $\mu = \sum_{i=1}^I \theta_i \mu^i$.

- A priori information on distribution of preferences of voters from group i is modeled as the inclusion $\mu^i \in M^i$, for some known subset $M^i \subset \Delta_L$ *which we assume to be nonempty convex compact set*.
- Output of particular interview is pair (i, j) , where $i \in \{1, \dots, I\}$ is selected at random according to probability distribution q , and j is the candidate preferred by respondent selected from group i at random, according to uniform distribution on the group.

Observation (outcome of an interview) *becomes*

$$\omega := (i, j) \in \Omega = \{1, \dots, I\} \times \{1, \dots, L\},$$
$$\text{Prob}\{\omega = (i, j)\} = p(i, j) := q_i \mu_j^i.$$

Hypotheses to decide upon are

$$H_\ell[q] : p \in \mathcal{P}_\ell[q] := \left\{ \{p_{ij} = q_i \mu_j^i\}_{\substack{1 \leq i \leq I, \\ 1 \leq j \leq L}} : \begin{array}{l} \mu^i \in M^i \forall i, \\ [\sum_i \theta_i \mu^i]_\ell \geq \delta + [\sum_i \theta_i \mu^i]_{\ell'} \forall (\ell' \neq \ell) \end{array} \right\}$$

$H_\ell[q]$, $\ell = 1, \dots, L$, states that the “signal” $\mu = [\mu^1; \dots; \mu^I]$ underlying distribution p of observations ω induces population-wide distribution $\sum_i \theta_i \mu^i$ of votes resulting in electing candidate ℓ with winning margin $\geq \delta$.

$$H_\ell[q] : p \in \mathcal{P}_\ell[q] := \left\{ \{p_{ij} = q_i \mu_j^i\}_{\substack{1 \leq i \leq I, \\ 1 \leq j \leq L}} : \begin{array}{l} \mu^i \in M^i \forall i, \\ [\sum_i \theta_i \mu^i]_\ell \geq \delta + [\sum_i \theta_i \mu^i]_{\ell'}, \forall (\ell' \neq \ell) \end{array} \right\}$$

Hypotheses $H_\ell[q]$ are of the form

$$H_\ell[q] = \{p = A[q]\boldsymbol{\mu} : \boldsymbol{\mu} := [\mu^1; \dots; \mu^L] \in \mathcal{M}^\ell, \\ [A[q]\boldsymbol{\mu}]_{ij} = q_i \mu_j^i,$$

where $\mathcal{M}^\ell$, $\ell = 1, \dots, L$, are nonempty nonintersecting convex compact subsets in $\underbrace{\Delta_L \times \dots \times \Delta_L}_I$

Note: Election Poll with K interviews corresponds to stationary K -repeated observation in Discrete o.s. with IL -element observation space Ω

$\Rightarrow$ Given K , we can design a near-optimal detector-based test $\mathcal{T}_K$ deciding via stationary K -repeated observation (i.e., via the outcomes of K interviews) on hypotheses $H_\ell[q]$, $\ell = 1, \dots, L$ up to trivial closeness

$$H_\ell[q] \text{ is close to } H_{\ell'}[q] \text{ iff } \ell = \ell'.$$

This test will predict the winner with reliability $1 - \text{Risk}[\mathcal{T}_K | H_1[q], \dots, H_L[q]]$.

$$H_\ell[q] = \{p = A[q]\mu : \mu := [\mu^1; \dots; \mu^L] \in \mathcal{M}^\ell\},$$

$$[A[q]\mu]_{ij} = q_i \mu_j^i,$$

Setting $\chi_{\ell\ell'} = \begin{cases} 0, & \ell = \ell' \\ 1, & \ell \neq \ell' \end{cases}$, we get

$$\text{Risk}[\mathcal{T}_K | H_1[q], \dots, H_L[q]] \leq \epsilon_K[q] := \left\| \left[\epsilon_{\ell\ell'}^K[q] \chi_{\ell\ell'} \right]_{\ell, \ell'=1}^L \right\|_{2,2},$$

where

$$\begin{aligned} \epsilon_{\ell\ell'}[q] &= \max_{\mu \in \mathcal{M}^\ell, \nu \in \mathcal{M}^{\ell'}} \sum_{i,j} \sqrt{[A[q]\mu]_{ij} [A[q]\nu]_{ij}} \\ &= \max_{\mu \in \mathcal{M}^\ell, \nu \in \mathcal{M}^{\ell'}} \underbrace{\sum_{i=1}^I q_i \left[\sum_{j=1}^L \sqrt{\mu_j^i \nu_j^i} \right]}_{\Phi(q; \mu, \nu)} \end{aligned}$$

Observe that $\Phi(q; \mu, \nu)$ is linear in q .

Let us carry out *Measurement Design* – optimize the value $\epsilon_K[q]$ in q .

Main observation: $\epsilon_K[q] = \Gamma(\Psi(q))$, where

- $\Gamma(Q) = \|[(Q_{\ell\ell'})^K \chi_{\ell\ell'}]_{\ell,\ell'=1}^L\|_{2,2}$ is efficiently computable *convex* and *entrywise nondecreasing* function on the space of nonnegative $L \times L$ matrices
- $\Psi(q)$ is matrix-valued function with efficiently computable *convex in q and nonnegative* entries

$$\Psi_{\ell\ell'}(q) = \max_{\mu \in \mathcal{M}^\ell, \nu \in \mathcal{M}^{\ell'}} \Phi(q; \mu, \nu)$$

$\Rightarrow$ *Optimal selection of q_i 's reduces to solving explicit convex problem*

$$\min_q \left\{ \Gamma(\Psi(q)) : q = [q_1; \dots; q_I] \geq 0, \sum_{i=1}^I q_i = 1 \right\}$$

How it Works: Measurement Design in Election Polls

Setup:

- Election Poll problem with L candidates and winning margin $\delta = 0.05$
- Reliability $\epsilon = 0.01$
- A priori information on voters' preferences in groups:

$$M^i = \{\mu^i \in \Delta_L : p_\ell^i - u_i \leq \mu_\ell^i \leq p_\ell^i + u_i, \ell \leq L\}$$

- p^i : randomly selected probabilistic vector
- u_i : uncertainty level

Sample of results:

L	I	Group sizes θ Uncertainty levels u	K_{ini}	q_{opt}	K_{opt}
2	2	$\theta = [0.50; 0.50]$ $u = [0.03; \mathbf{1.00}]$	1212	$[0.44; 0.56]$	1194
2	2	$[0.50; 0.50]$ $[0.02; \mathbf{1.00}]$	$\mathbf{2699}$	$[0.00; 1.00]$	$\mathbf{1948}$
3	3	$[0.33; 0.33; 0.33]$ $[0.02; 0.03; \mathbf{1.00}]$	$\mathbf{3177}$	$[0.00; 0.46; 0.54]$	$\mathbf{2726}$
5	4	$[0.25; 0.25; 0.25; 0.25]$ $[0.02; 0.02; 0.03; \mathbf{1.00}]$	$\mathbf{2556}$	$[0.00; 0.13; 0.32; 0.55]$	$\mathbf{2086}$
5	4	$[0.25; 0.25; 0.25; 0.25]$ $[\mathbf{1.00}; \mathbf{1.00}; \mathbf{1.00}; \mathbf{1.00}]$	4788	$[0.25; 0.25; 0.25; 0.25]$	4788

Effect of measurement design. K_{ini} and K_{opt} are the poll sizes required for 0.99-reliable prediction of the winner when $q_i = \theta_i$ and $q = q_{\text{opt}}$, respectively.

Note: Uncertainty = $\mathbf{1.00} \Leftrightarrow$ No a priori information

- *In numerous situations, we do have partial control of observation scheme and thus can look for optimal Measurement Design.*
- However, situations where optimal Measurement Design can be found efficiently, like in design of Election Polls, are rare.
- Two other examples of such situation are *Poisson o.s. and Gaussian o.s. with time control.*

Poisson o.s. with time control.

Typical models where Poisson o.s. arises are as follows:

- there exists a “signal” x known to belong to some convex compact set $\subset \mathbb{R}^n$

For example, in **P**ositron **E**mission **T**omography, x is (discretized) density of radioactive tracer administered to patient

- We observe random vector $\omega \in \mathbb{R}^m$ with independent entries $\omega_i \sim \text{Poisson}(a_i^T x)$, and want to make inferences on x .

In PET, patient is injected radioactive tracer which concentrates in areas we are interested in. Tracer disintegrates, and every disintegration act results in pair of gamma-quants flying in opposite directions along a randomly oriented line passing through disintegration point. This line is registered when two detector cells are (nearly) simultaneously hit:

The data acquired in PET study are the numbers ω_i of lines registered in *bins* (pairs of detector cells) $i = 1, \dots, m$ over a time horizon T . We have

$$\omega_i \sim \text{Poisson}(T \sum_{j=1}^n p_{ij} x_j)$$

$$\left[\begin{array}{l} p_{ij}: \text{probability for line emanating from voxel } j = 1, \dots, n \\ \text{to cross pair } i = 1, \dots, m \text{ of detector cells} \end{array} \right]$$

$$\Rightarrow A = T \left[p_{ij} \right]_{i \leq m, j \leq n}$$

In some situations, the sensing matrix A can be partially controlled:

$$A = A[q] := \text{Diag}\{q\} A_*$$

- A_* : given $m \times n$ matrix;
- $q \in \mathcal{Q}$: vector of control parameters.

For instance, in a whole body PET scan the position of the patient w.r.t. the scanner is updated several times to cover the entire body.

The data acquired in position ι form subvector ω^ι in the entire observation $\omega = [\omega^1; \dots; \omega^I]$:

$$\omega_i^\iota \sim \text{Poisson}([t_\iota A^\iota x]_i), \quad 1 \leq i \leq \bar{m} = m/I$$

$$\left[\begin{array}{l} A^\iota : \text{ given matrices; } t_\iota : \text{ duration of study in position } \iota \end{array} \right]$$

implying that $A = \text{Diag}\{q\}A_*$ with properly selected A_* and q of the form

$$q = \left[\underbrace{t_1; \dots; t_1}_{\bar{m}}; \dots; \underbrace{t_I; \dots; t_I}_{\bar{m}} \right]$$

$H_\ell^q : \omega_i \sim \text{Poisson}([A[q]x]_i)$ are independent across $i \leq m$ and $x \in X_\ell$
 $A = A[q] := \text{Diag}\{q\}A_*$

- A_* : given $m \times n$ matrix; • $q \in \mathcal{Q}$: control parameters.

Our goal is to decide, up to a given closeness $\mathcal{C}$, on L hypotheses on the distribution of Poisson observation ω :

$H_\ell^q : \omega \sim \text{Poisson}([A[q]x]_1) \times \dots \times \text{Poisson}([A[q]x]_m)$
 $x \in X_\ell, X_\ell$: given convex compact sets, $1 \leq \ell \leq L$.

- By our theory, the (upper bound on the) $\mathcal{C}$ -risk of near-optimal test deciding on H_ℓ^q , $\ell = 1, \dots, L$, is

$$\epsilon(q) = \left\| [\exp\{\text{Opt}_{\ell\ell'}(q)\} \chi_{\ell\ell'}]_{\ell,\ell'=1}^L \right\|_{2,2}$$

$$\left[\chi_{\ell\ell'} = \begin{cases} 0, & (\ell, \ell') \in \mathcal{C} \\ 1, & (\ell, \ell') \notin \mathcal{C} \end{cases}, \text{Opt}_{\ell\ell'}(q) = \max_{u \in X_\ell, v \in X_{\ell'}} -\frac{1}{2} \sum_{i=1}^m \left(\sqrt{[A[q]u]_i} - \sqrt{[A[q]v]_i} \right)^2 \right]$$

Similarly to the Election Polls, $\epsilon(q) = \Gamma(\Psi(q))$, where

- $\Gamma(Q) = \left\| [\exp\{Q_{\ell\ell'}\} \chi_{\ell\ell'}]_{\ell,\ell'=1}^L \right\|_{2,2}$ is a convex entrywise nondecreasing function of $L \times L$ matrix Q

- $[\Psi(q)]_{\ell\ell'} = \exp \left\{ \max_{u \in X_\ell, v \in X_{\ell'}} \sum_{i=1}^m q_i \left(\sqrt{[A_*u]_i [A_*v]_i} - \frac{1}{2}[A_*u]_i - \frac{1}{2}[A_*v]_i \right) \right\}$

Note: same as in the Election Polls case, $[\Psi(q)]_{\ell\ell'}$ is efficiently computable and convex in q

$\Rightarrow$ Assuming the set $\mathcal{Q} \subset \mathbb{R}_+^m$ of allowed controls q is convex, optimizing $\epsilon(q)$ over $q \in \mathcal{Q}$ is an *explicit convex optimization problem*.

Measurement Design problem in Gaussian o.s.

$$\omega = A[q]x + \xi, \xi \sim \mathcal{N}(0, I_m)$$

[• $A[q]$ partially controlled sensing matrix; • $q \in \mathcal{Q}$: control parameters.]

is efficiently solvable if

$$A[q] = \text{Diag}\{\sqrt{q_1}, \dots, \sqrt{q_m}\} A_*,$$

$$\mathcal{Q} \subset \mathbb{R}_+^m \text{ is a convex compact set}$$

In this case, minimizing $\mathcal{Q}$ -risk of test deciding up to closeness $\mathcal{C}$ on L hypotheses

$$H_\ell^q : \omega \sim \mathcal{N}(A[q]x, I_m), \quad x \in X_\ell, \quad 1 \leq \ell \leq L$$

associated with nonempty convex compact sets X_ℓ reduces to solving convex problem

$$\min_{q \in \mathcal{Q}} \Gamma(\Psi(q))$$

where

$$\Gamma(Q) = \| [\exp\{Q_{\ell\ell'}/8\} \chi_{\ell\ell'}]_{\ell, \ell' \leq L} \|_{2,2}$$

is convex entrywise nondecreasing function of $L \times L$ matrix Q , and

$$[\Psi(q)]_{\ell\ell'} = \max_{u \in X_\ell, v \in X_{\ell'}} [-\|A[q](u - v)\|_2^2] = - \min_{u \in X_\ell, v \in X_{\ell'}} (u - v)^T A_*^T \text{Diag}\{q\} A_* (u - v)$$

is efficiently computable convex function of $q \in \mathcal{Q}$.

Illustration: In some applications, “the physics” beyond Gaussian o.s. $\omega = Ax + \xi$ is as follows.

There are m sensors measuring analogous vector-valued continuous time signal x (nearly constant on the observation horizon) in the presence of noise. The output of sensor $\#i$ is

$$\omega_i = \frac{1}{|\Delta_i|} \int_{\Delta_i} [a_{i,*}^T x + B_i(t)] dt$$

$$\left[\begin{array}{l} \bullet \Delta_i : \text{continuous time interval} \\ \bullet B_i(t) : \text{“Brownian motion:” } \frac{1}{|\Delta|} \int_{\Delta} B_i(t) dt \sim \mathcal{N}(0, |\Delta|^{-1}), \\ \int_{\Delta} B_i(t) dt, \int_{\Delta'} B_i(t) dt \text{ are independent when } \Delta \cap \Delta' = \emptyset \\ \bullet \text{Brownian motions } B_i(t) \text{ are independent across } i \end{array} \right]$$

- When all sensors work in parallel for unit time, we arrive at the standard Gaussian o.s. $\omega = A_* x + \xi, \xi \sim \mathcal{N}(0, I_m)$.
- When sensors work on consecutive segments $\Delta_1, \dots, \Delta_m$ of durations $q_i = |\Delta_i|$, we arrive at

$$\omega_i = a_{i,*}^T x + q_i^{-1/2} \xi_i, \xi_i \sim \mathcal{N}(0, 1) \text{ are independent across } i$$

Rescaling observations according to $\omega_i \mapsto \sqrt{q_i} \omega_i$, we arrive at the desired partially controlled observation scheme

$$o\omega = \text{Diag}\{\sqrt{q_1}, \dots, \sqrt{q_m}\} A_* x + \xi, \xi \sim \mathcal{N}(0, I_m)$$

A natural selection of $\mathcal{Q}$ is, e.g., $\mathcal{Q} = \{q \geq 0 : \sum_i q_i = m\}$ (setting “time budget” to the same value as in the case of consecutive observations of duration 1 each).

HYPOTHESIS TESTING, III

- *Beyond simple observation schemes*
 - *Simple families of distributions*
 - *What is affine? Quadratic lifting*

Observation: A “common denominator” of minimum risk detectors for simple o.s.’s is their *affinity* in observations:

- the optimal detectors in Gaussian and Poisson o.s.’s are affine “as they are”
- encoding observation space $\Omega = \{1, \dots, d\}$ of Discrete o.s. by vertices $e_i, i = 1, \dots, d$, of the standard simplex $\Delta_d = \{x \in \mathbb{R}^d : x \geq 0, \sum_j x_j = 1\}$, every function on Ω becomes affine
 $\Rightarrow$ *we can treat optimal detector in Discrete o.s. as affine function on $\mathbb{R}^d$.*
- operations with optimal detectors induced by taking direct products of basic simple o.s.’s or passing to repeated observations preserve affinity.

Claim: “Reasonable” (perhaps, sub-optimal) affine detectors can be found, in a computationally efficient way, in many important situations which are beyond simple o.s.’s.

$$\begin{aligned}\mathcal{O} &= ((\Omega, \Pi), \{p_\mu(\cdot) : \mu \in \mathcal{M}\}, \mathcal{F}) \\ \Phi_{\mathcal{O}}(\phi, \mu) &= \ln \left(\int \exp\{\phi(\omega)\} p_\mu(\omega) \Pi(d\omega) \right) : \mathcal{F} \times \mathcal{M} \rightarrow \mathbb{R} \\ &\text{convex in } \phi, \text{ concave in } \mu\end{aligned}$$

In a simple o.s., our strategy to build a good detector for two hypotheses associated with convex compact sets $M_1 \subset \mathcal{M}$, $M_2 \subset \mathcal{M}$ was as follows:

In order to find minimum risk detector – to solve the optimization problem

$$\min_{\phi \in \mathcal{F}} \max \left[\max_{\mu \in M_1} \int \exp\{-\phi(\omega)\} p_\mu(\omega) \Pi(d\omega), \max_{\nu \in M_2} \int \exp\{\phi(\omega)\} p_\nu(\omega) \Pi(d\omega) \right]$$

reduced to

$$\begin{aligned}\min_{\phi \in \mathcal{F}} \frac{1}{2} \left[\max_{\mu \in M_1} \ln \left(\int \exp\{-\phi(\omega)\} p_\mu(\omega) \Pi(d\omega) \right) + \max_{\nu \in M_2} \ln \left(\int \exp\{\phi(\omega)\} p_\nu(\omega) \Pi(d\omega) \right) \right] \\ \Leftrightarrow \min_{\phi \in \mathcal{F}} \frac{1}{2} \left[\max_{\mu \in M_1} \Phi_{\mathcal{O}}(-\phi, \mu) + \max_{\nu \in M_2} \Phi_{\mathcal{O}}(\phi, \nu) \right]\end{aligned}$$

- We intend to use the same scheme, but replace “ $\phi \in \mathcal{F}$ ” with “*affine* ϕ ”, and replace $\Phi_{\mathcal{O}}$ with *convex in ϕ and concave in μ upper bound on*

$$\ln \left(\int \exp\{\phi(\omega)\} p_\mu(\omega) \Pi(d\omega) \right).$$

Setup

- Given an observation space $\Omega = \mathbb{R}^d$, consider a triple $\mathcal{H}, \mathcal{M}, \Phi$, where
 - $\mathcal{H}$ is a nonempty closed convex set in Ω symmetric w.r.t. the origin,
 - $\mathcal{M}$ is a compact convex set in some $\mathbb{R}^n$,
 - $\Phi(h; \mu) : \mathcal{H} \times \mathcal{M} \rightarrow \mathbb{R}$ is a continuous function *convex in $h \in \mathcal{H}$* and *concave in $\mu \in \mathcal{M}$* .
- $\mathcal{H}, \mathcal{M}$, and Φ specify a *family $\mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$ of probability distributions* on Ω . A probability distribution P belongs to the family iff there exists $\mu \in \mathcal{M}$ such that

$$\ln \left(\int_{\Omega} e^{h^T \omega} P(d\omega) \right) \leq \Phi(h; \mu) \quad \forall h \in \mathcal{H} \quad (*)$$

We refer to μ ensuring $(*)$ as to *parameter* of distribution P .

Warning: A distribution P may have many different parameters!

- We refer to triple $\mathcal{H}, \mathcal{M}, \Phi$ satisfying the above requirements as to *regular data*, and to $\mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$ – as to the *simple family of distributions* induced by these data.

Example 1: Gaussian and sub-Gaussian distributions.

When $\mathcal{M} = \{(u, \Theta)\} \subset \mathbb{R}^d \times \text{int } \mathbf{S}_+^d$ is a convex compact set such that $\Theta \succ 0$ for all $(u, \Theta) \in \mathcal{M}$, $\mathcal{H} = \mathbb{R}^d$ and

$$\Phi(h; u, \Theta) = h^T y + \frac{1}{2} h^T \Theta h, \quad \mathcal{S} = \mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$$

contains all probability distributions P which are sub-Gaussian with parameters (u, Θ) , meaning that

$$\ln \left(\int_{\Omega} e^{h^T \omega} P(d\omega) \right) \leq h^T u + \frac{1}{2} h^T \Theta h \quad \forall h, \quad (1)$$

and, in addition, the “parameter” (u, Θ) belongs to $\mathcal{M}$.

Note:

- $\mathcal{N}(u, \Theta) \in \mathcal{S}$ whenever $(u, \Theta) \in \mathcal{M}$; for $P = \mathcal{N}(u, \Theta)$, (1) is an identity
- Whenever P is sub-Gaussian with parameters (u, Θ) , u is the expectation of P

Example 2: Poisson distributions.

When $\mathcal{M} \subset \mathbb{R}_+^d$ is a convex compact set, $\mathcal{H} = \mathbb{R}^d$ and

$$\Phi(h; \mu) = \sum_{i=1}^d \mu_i (e^{h_i} - 1),$$

$\mathcal{S} = \mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$ contains distributions of all d -dimensional random vectors ω_i with independent across i entries

$$\omega_i \sim \text{Poisson}(\mu_i)$$

such that $\mu = [\mu_1; \dots; \mu_d] \in \mathcal{M}$.

Example 3: Discrete distributions.

When

$$\mathcal{M} = \{\mu \in \mathbb{R}^d : \mu \geq 0, \sum_j \mu_j = 1\}$$

is the probabilistic simplex in $\mathbb{R}^d$, $\mathcal{H} = \mathbb{R}^d$ and

$$\Phi(h; \mu) = \ln \left(\sum_{i=1}^d \mu_i e^{h_i} \right),$$

$\mathcal{S} = \mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$ contains all discrete distributions supported on the vertices of the probabilistic simplex.

Example 4: Distributions with bounded support.

Let $X \subset \mathbb{R}^d$ be a nonempty convex compact set with support function $\phi_X(\cdot)$:

$$\phi_X(y) = \max_{x \in X} y^T x : \mathbb{R}^d \rightarrow \mathbb{R}.$$

When $\mathcal{M} = X$, $\mathcal{H} = \mathbb{R}^d$ and

$$\Phi(h; \mu) = h^T \mu + \frac{1}{8} [\phi_X(h) + \phi_X(-h)]^2, \quad (2)$$

$\mathcal{S} = \mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$ contains all probability distributions supported on X , and for such a distribution P , $\mu = \int_X \omega P(d\omega)$ is a parameter of P .

Note: When G , $0 \in G$, is a convex compact set, the conclusion in Example 4 remains valid when function (2) is replaced with the smaller function

$$\Phi(h; \mu) = \min_{g \in G} \left[\mu^T (h - g) + \frac{1}{8} [\phi_X(h - g) + \phi_X(g - h)]^2 + \phi_X(g) \right].$$

“Calculus” of simple families of probability distributions

- [summation] For $1 \leq \ell \leq L$, let λ_ℓ be reals, and let $\mathcal{H}_\ell, \mathcal{M}_\ell, \Phi_\ell$ be regular data with common observation space: $\mathcal{H}_\ell \subset \Omega = \mathbb{R}^d$. Setting

$$\begin{aligned}\mathcal{H} &= \{h \in \mathbb{R}^d : \lambda_\ell h \in \mathcal{H}_\ell, 1 \leq \ell \leq L\}, \\ \mathcal{M} &= \mathcal{M}_1 \times \dots \times \mathcal{M}_L, \\ \Phi(h; \mu_1, \dots, \mu_L) &= \sum_{\ell=1}^L \Phi_\ell(\lambda_\ell h; \mu_\ell),\end{aligned}$$

we get regular data with the following property:

Whenever random vectors $\xi_\ell \sim P_\ell \in \mathcal{S}[\mathcal{H}_\ell, \mathcal{M}_\ell, \Phi_\ell]$, $1 \leq \ell \leq L$, are independent across ℓ , the distribution P of the random vector

$$\xi = \sum_{\ell=1}^L \lambda_\ell \xi_\ell$$

belongs to $\mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$. Denoting by μ_ℓ parameters of P_ℓ , $\mu = [\mu_1; \dots; \mu_L]$ can be taken as parameter of P .

- **[direct product]** For $1 \leq \ell \leq L$, let $\mathcal{H}_\ell, \mathcal{M}_\ell, \Phi_\ell$ be regular data with observation spaces $\Omega_\ell = \mathbb{R}^{d_\ell}$. Setting

$$\begin{aligned}\mathcal{H} &= \mathcal{H}_1 \times \dots \times \mathcal{H}_L \subset \Omega = \mathbb{R}^{d_1 + \dots + d_L}, \\ \mathcal{M} &= \mathcal{M}_1 \times \dots \times \mathcal{M}_L, \\ \Phi(h_1, \dots, h_L; \mu_1, \dots, \mu_L) &= \sum_{\ell=1}^L \Phi_\ell(h_\ell; \mu_\ell),\end{aligned}$$

we get regular data with the following property:

Whenever $P_\ell \in \mathcal{S}[\mathcal{H}_\ell, \mathcal{M}_\ell, \Phi_\ell]$, $1 \leq \ell \leq L$, the direct product distribution $P = P_1 \times \dots \times P_L$ belongs to $\mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$. Denoting by μ_ℓ parameters of P_ℓ , $\mu = [\mu_1; \dots; \mu_L]$ can be taken as parameter of P .

- [marginal] Let $\mathcal{H}, \mathcal{M}, \Phi$ be regular data with observation space $\mathbb{R}^d$, and let

$$\omega \mapsto A\omega + a : \mathbb{R}^d \mapsto \Omega = \mathbb{R}^\delta.$$

Setting

$$\bar{\mathcal{H}} = \{h \in \mathbb{R}^\delta : A^T h \in \mathcal{H}\}, \bar{\Phi}(h; \mu) = h^T a + \Phi(A^T h; \mu),$$

we get regular data $\bar{\mathcal{H}}, \mathcal{M}, \bar{\Phi}$ with the following property:

Whenever $\xi \sim P \in \mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi]$, the distribution $\bar{P}$ of the random variable $\omega = A\xi + a$ belongs to the simple family $\mathcal{S}[\bar{\mathcal{H}}, \mathcal{M}, \bar{\Phi}]$, and parameter of P is a parameter of $\bar{P}$ as well.

Main observation: When deciding on simple families of distributions, affine tests and their risks can be efficiently computed via Convex Programming:

Theorem. Let $\mathcal{H}_\chi, \mathcal{M}_\chi, \Phi_\chi$, $\chi = 1, 2$, be two collections of regular data with compact $\mathcal{M}_1, \mathcal{M}_2$ and $\mathcal{H}_1 = \mathcal{H}_2 =: \mathcal{H}$, and let

$$\Psi(h) = \max_{\mu_1 \in \mathcal{M}_1, \mu_2 \in \mathcal{M}_2} \underbrace{\frac{1}{2} [\Phi_1(-h; \mu_1) + \Phi_2(h; \mu_2)]}_{\Phi(h; \mu_1, \mu_2)} : \mathcal{H} \rightarrow \mathbb{R}$$

Then Ψ is efficiently computable continuous convex function, and for $h \in \mathcal{H}$, setting

$$\phi_h(\omega) = h^T \omega + \underbrace{\frac{1}{2} \left[\max_{\mu_1 \in \mathcal{M}_1} \Phi_1(-h; \mu_1) - \max_{\mu_2 \in \mathcal{M}_2} \Phi_2(h; \mu_2) \right]}_{\kappa},$$

one has

$$\text{Risk}[\phi | \mathcal{P}_1, \mathcal{P}_2] \leq \exp\{\Psi(h)\}, \text{ where } \mathcal{P}_\chi = \mathcal{S}[\mathcal{H}, \mathcal{M}_\chi, \Phi_\chi], \chi = 1, 2.$$

In particular, if convex-concave function $\Phi(h; \mu_1, \mu_2)$ possesses a saddle point $h_*, (\mu_1^*, \mu_2^*)$ on $\mathcal{H} \times (\mathcal{M}_1 \times \mathcal{M}_2)$, the affine detector

$$\phi_*(\omega) = h_*^T \omega + \frac{1}{2} [\Phi_1(-h_*; \mu_1^*) - \Phi_2(h_*; \mu_2^*)]$$

admits risk bound

$$\text{Risk}[\phi | \mathcal{P}_1, \mathcal{P}_2] \leq \exp\{\Phi(h_*; \mu_1^*, \mu_2^*)\}.$$

Proof. Indeed, let $h \in \mathcal{H}$, and let

$$\mu_1^* \in \underset{\mu_1 \in \mathcal{M}_1}{\text{Argmax}} \Phi_1(-h; \mu_1), \mu_2^* \in \underset{\mu_2 \in \mathcal{M}_2}{\text{Argmax}} \Phi_2(h; \mu_2).$$

When $P \in \mathcal{P}_1 := \mathcal{S}[\mathcal{H}, \mathcal{M}_1, \Phi_1]$, there is $\mu_1 \in \mathcal{M}_1$ such that

$$\mathbf{E}_{\omega \sim P} \left\{ e^{-h^T \omega} P(d\omega) \right\} \leq e^{\Phi_1(-h; \mu_1)},$$

so that

$$\mathbf{E}_{\omega \sim P} \left\{ e^{-\phi(\omega)} P(d\omega) \right\} \leq e^{\Phi_1(-h; \mu_1^*) - \kappa} = e^{\Psi(h)},$$

implying that

$$\text{Risk}_1[\phi | \mathcal{P}_1, \mathcal{P}_2] \leq e^{\Psi(h)}.$$

Similarly, when $\mu_1 \in \mathcal{M}_1$, for some $\mu_2 \in \mathcal{M}_2$ one has

$$\mathbf{E}_{\omega \sim P} \left\{ e^{h^T \omega} P(d\omega) \right\} \leq e^{\Phi_2(h; \mu_2)},$$

$$\mathbf{E}_{\omega \sim P} \left\{ e^{\phi(\omega)} P(d\omega) \right\} \leq e^{\Phi_2(h; \mu_2^*) + \kappa} = e^{\Psi(h)},$$

and

$$\text{Risk}_2[\phi | \mathcal{P}_1, \mathcal{P}_2] \leq e^{\Psi(h)}.$$

□

Numerical illustration. Given convex compact subsets X_1 and X_2 of X , and observation

$$\omega = Ax + \sigma A \text{Diag} \{ \sqrt{x_1}, \dots, \sqrt{x_n} \} \xi \quad [\xi \sim \mathcal{N}(0, I_n)]$$

of an unknown signal x known to belong to a given convex compact set $X \subset \mathbb{R}_{++}^n$, we want to decide on the hypotheses

$$H_\chi : x \in X_\chi, \quad \chi = 1, 2$$

with risk ≤ 0.01 .

Novelty: *Noise intensity depends on the signal!*

- Introducing regular data $\mathcal{H}_\chi = \mathbb{R}^n$, $\mathcal{M}_\chi = X_\chi$,

$$\Phi_\chi(h, \mu) = h^T A \mu + \frac{1}{2} \sigma^2 h^T [A \text{Diag} \{ \mu \} A^T] h \quad [\chi = 1, 2]$$

distribution of observations under H_χ belongs to $\mathcal{S}[\mathcal{H}, \mathcal{M}_\chi, \Phi_\chi]$.

- An affine detector for families $\mathcal{P}_\chi$ of distributions obeying H_χ , $\chi = 1, 2$, is given by the saddle point of the function

$$\Phi(h; \mu_1, \mu_2) := \frac{1}{2} \left[h^T [\mu_2 - \mu_1] + \frac{1}{2} \sigma^2 h^T A \text{Diag} \{ \mu_1 + \mu_2 \} A^T h \right]$$

Data: $n = 16$, $\sigma = 0.1$, target risk 0.01,

- $A = U \text{Diag} \{0.01^{(i-1)/15}, i \leq 16\} V^T$ with random orthogonal U, V ,

- $X_1 = \left\{ x \in \mathbb{R}^{16} : \begin{array}{l} 0.001 \leq x_1 \leq \delta \\ 0.001 \leq x_i \leq 1, i \geq 2 \end{array} \right\},$

- $X_2 = \left\{ x \in \mathbb{R}^{16} : \begin{array}{l} 2\delta \leq x_1 \leq 1 \\ 0.001 \leq x_i \leq 1, i \geq 2 \end{array} \right\}$

Results:

$\delta = 0.1 \Rightarrow \text{Risk}[\phi_* | \mathcal{P}_1, \mathcal{P}_2] = 0.4346 \Rightarrow$ 6-repeated observation

$\delta = 0.01 \Rightarrow \text{Risk}[\phi_* | \mathcal{P}_1, \mathcal{P}_2] = 0.9201 \Rightarrow$ 56-repeated observation

- Safe “Gaussian o.s. approximation” of the above observation scheme requires 37-repeated observations to handle $\delta = 0.1$ and 3685-repeated observation to handle $\delta = 0.01$.

Sub-Gaussian case. For $\chi = 1, 2$, let $U_\chi \subset \Omega = \mathbb{R}^d$ and $\mathcal{V}_\chi \subset \text{int } S_+^d$ be convex compact sets. Setting

$$\mathcal{M}_\chi = U_\chi \times \mathcal{V}_\chi, \quad \Phi(h; u, \Theta) = h^T u + \frac{1}{2} h^T \Theta h : \mathcal{H} \times \mathcal{M}_\chi \rightarrow \mathbb{R},$$

the regular data $\mathcal{H} = \mathbb{R}^d$, $\mathcal{M}_\chi$, Φ specify the families $\mathcal{P}_\chi = \mathcal{S}[\mathbb{R}^d, U_\chi \times \mathcal{V}_\chi, \Phi]$ of sub-Gaussian distributions with parameters from $U_\chi \times \mathcal{V}_\chi$.

Saddle point problem responsible for design of affine detector for $\mathcal{P}_1, \mathcal{P}_2$ reads

$$\text{SadVal} = \min_{h \in \mathbb{R}^d} \max_{\substack{u_1 \in U_1, u_2 \in U_2 \\ \Theta_1 \in \mathcal{V}_1, \Theta_2 \in \mathcal{V}_2}} \frac{1}{2} [h^T (u_2 - u_1) + \frac{1}{2} h^T [\Theta_1 + \Theta_2] h]$$

- Saddle point $(h_*; (u_1^*, u_2^*, \Theta_1^*, \Theta_2^*))$ does exist and satisfies

$$h_* = [\Theta_1^* + \Theta_2^*]^{-1} [u_1^* - u_2^*],$$

$$\text{SadVal} = -\frac{1}{4} [u_1^* - u_2^*] [\Theta_1^* + \Theta_2^*]^{-1} [u_1^* - u_2^*] = -\frac{1}{4} h_*^T [u_1^* - u_2^*]$$

- The associated affine detector and its risk are given by

$$\phi_*(\omega) = h_*^T [\omega - \frac{1}{2} [u_1^* + u_2^*]] = [u_1^* - u_2^*]^T [\Theta_1^* + \Theta_2^*]^{-1} [\omega - \frac{1}{2} [u_1^* + u_2^*]]$$

$$\text{Risk}[\phi_* | \mathcal{P}_1, \mathcal{P}_2] \leq \exp\{\text{SadVal}\} = \exp\{-\frac{1}{4} [u_1^* - u_2^*] [\Theta_1^* + \Theta_2^*]^{-1} [u_1^* - u_2^*]\}$$

Note: In the *symmetric case* $\mathcal{V}_1 = \mathcal{V}_2$ $(h_*; (u_1^*, u_2^*, \Theta_1^*, \Theta_2^*))$ can be selected to have $\Theta_1^* = \Theta_2^* =: \Theta_*$. In this case, *the affine detector we end up with is the minimum risk detector for $\mathcal{P}_1, \mathcal{P}_2$.*

What is “affine?” Quadratic Lifting

We have developed a technique for building reasonable *affine* detectors for simple families of distributions.

Given observation $\zeta \sim P$, we can subject it to *nonlinear* transformation $\zeta \mapsto \omega = \psi(\zeta)$, e.g., to *quadratic lifting*

$$\zeta \mapsto \omega = (\zeta, \zeta\zeta^T)$$

and treat as our observation ω rather than ζ . *Affine in ω detectors are nonlinear in ζ .*

Example: Detectors affine in the quadratic lifting $\omega = (\zeta, \zeta\zeta^T)$ of ζ are exactly the *quadratic* functions of ζ .

- We can try to apply our machinery for building affine detectors to nonlinear transformations of true observations, thus arriving at nonlinear detectors.

Bottleneck: To apply the outlined strategy to a pair $\mathcal{P}_1, \mathcal{P}_2$ of families of distributions of interest, *we need to cover the families $\mathcal{P}_1^+, \mathcal{P}_2^+$ of distributions of $\omega = \psi(\zeta)$ induced by distributions $P \in \mathcal{P}_\chi$ of ζ , $\chi = 1, 2$, by simple families of distributions.*

What is ahead: Simple “coverings” of quadratic lifts of (sub)Gaussian distributions.

Situation: Given are:

- a compact nonempty set $U \subset \mathbb{R}^n$
- an affine mapping $u \mapsto \mathcal{A}(u) = A[u; 1] : \mathbb{R}^n \rightarrow \mathbb{R}^d$
- a convex compact set $\mathcal{V} \subset \text{int } \mathbf{S}_+^d$.
- The above data specify families of probability distributions of random observations

$$\omega = (\zeta, \zeta \zeta^T), \quad \zeta = \mathcal{A}(u) + \xi \in \mathbb{R}^d, \quad (*)$$

specifically,

- the family $\mathcal{G}[U, \mathcal{V}]$ of all distributions of ω induced by deterministic $u \in U$ and *Gaussian* noise $\xi \sim \mathcal{N}(0, \Theta \in \mathcal{V})$
- the family $\mathcal{SG}[U, \mathcal{V}]$ of all distributions of ω induced by deterministic $u \in U$ and *sub-Gaussian*, with parameters $(0, \Theta \in \mathcal{V})$ noise ξ

Goal: To cover $\mathcal{G}$ (or $\mathcal{SG}$) by a simple family of distributions.

Gaussian case

Proposition. *Given the above data $U, \mathcal{A}(u) = A[u; 1], \mathcal{V}$, let us select*

- $\gamma \in (0, 1)$
- *a computationally tractable convex compact set*

$$\mathcal{Z} \subset \mathcal{Z}^+ = \{Z \in \mathbf{S}^{n+1} : Z \succeq 0, Z_{n+1,n+1} = 1\}$$

such that $[u; 1][u; 1]^T \in \mathcal{Z} \ \forall u \in U$

- *A matrix $\Theta_* \in \mathbf{S}^d$ and $\delta \in [0, 2]$ such that*

$$\forall (\Theta \in \mathcal{V}) : \Theta \preceq \Theta_* \quad \text{and} \quad \|\Theta^{1/2}\Theta_*^{-1/2} - I_d\|_{2,2} \leq \delta$$

Let us set

$$B = \begin{bmatrix} A \\ 0, \dots, 0, 1 \end{bmatrix} \in \mathbb{R}^{(d+1) \times (n+1)}, \quad \mathcal{M} = \mathcal{V} \times \mathcal{Z}$$

$$\mathcal{H} = \{(h, H) \in \mathbb{R}^d \times \mathbf{S}^d : -\gamma\Theta_*^{-1} \preceq H \preceq \gamma\Theta_*^{-1}\}$$

$$\begin{aligned} \Phi_{\mathcal{A}, \mathcal{Z}}(h, H; \Theta, Z) = & -\frac{1}{2} \ln \text{Det}(I - \Theta_*^{1/2} H \Theta_*^{1/2}) + \frac{1}{2} \text{Tr}([\Theta - \Theta_*]H) + \frac{\delta(2+\delta)\|\Theta_*^{1/2} H \Theta_*^{1/2}\|_F^2}{2(1-\|\Theta_*^{1/2} H \Theta_*^{1/2}\|_F)} \\ & \quad \quad \quad [\|\cdot\|_F - \text{Frobenius norm}] \\ & + \frac{1}{2} \text{Tr} \left(Z B^T \left[\left[\begin{array}{c|c} H & h \\ \hline h^T & 1 \end{array} \right] + [H, h]^T [\Theta_*^{-1} - H]^{-1} [H, h] \right] B \right) : \mathcal{H} \times \mathcal{M} \rightarrow \mathbb{R} \end{aligned}$$

Then $\mathcal{H}, \mathcal{M}, \Phi_{\mathcal{A}, \mathcal{Z}}$ is efficiently computable regular data, and

$$\mathcal{G}[U, \mathcal{V}] \subset \mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi_{\mathcal{A}, \mathcal{Z}}].$$

Sub-Gaussian case

Proposition. *Given the above data U , $\mathcal{A}(u) = A[u; 1]$, $\mathcal{V}$, let us select*

- $\gamma, \gamma^+ \in (0, 1)$ with $\gamma < \gamma^+$
- a computationally tractable convex compact set

$$\mathcal{Z} \subset \mathcal{Z}^+ = \{Z \in \mathbf{S}^{n+1} : Z \succeq 0, Z_{n+1,n+1} = 1\}$$

such that $[u; 1][u; 1]^T \in \mathcal{Z} \quad \forall u \in U$

- A matrix $\Theta_* \in \mathbf{S}^d$ and $\delta \in [0, 2]$ such that

$$\forall (\Theta \in \mathcal{V}) : \Theta \preceq \Theta_* \quad \text{and} \quad \|\Theta^{1/2}\Theta_*^{-1/2} - I_d\|_{2,2} \leq \delta$$

Let us set

$$B = \begin{bmatrix} A \\ 0, \dots, 0, 1 \end{bmatrix} \in \mathbb{R}^{(d+1) \times (n+1)},$$

$$\mathcal{H} = \{(h, H) \in \mathbb{R}^d \times \mathbf{S}^d : -\gamma\Theta_*^{-1} \preceq H \preceq \gamma\Theta_*^{-1}\}, \quad \mathcal{M} = \mathcal{Z}$$

$$\mathcal{H}^+ = \{(h, H, G) \in \mathbb{R}^d \times \mathbf{S}^d \times \mathbf{S}^d : -\gamma^+\Theta_*^{-1} \preceq H \preceq G \preceq \gamma^+\Theta_*^{-1}, 0 \preceq G\}$$

$$\Phi_{\mathcal{A}, \mathcal{Z}}(h, H; Z) = \min_{G: (h, H, G) \in \mathcal{H}^+} \left\{ -\frac{1}{2} \ln \text{Det}(I - \Theta_*^{1/2} G \Theta_*^{1/2}) \right. \\ \left. + \frac{1}{2} \text{Tr} \left(Z B^T \left[\begin{array}{c|c} H & h \\ \hline h^T & 1 \end{array} \right] + [H, h]^T [\Theta_*^{-1} - G]^{-1} [H, h] \right) B \right\} : \mathcal{H} \times \mathcal{M} \rightarrow \mathbb{R}$$

Then $\mathcal{H}, \mathcal{M}, \Phi_{\mathcal{A}, \mathcal{Z}}$ is efficiently computable regular data, and

$$\mathcal{SG}[U, \mathcal{V}] \subset \mathcal{S}[\mathcal{H}, \mathcal{M}, \Phi_{\mathcal{A}, \mathcal{Z}}].$$

How to specify $\mathcal{Z}$? To apply the above construction, one should specify a computationally tractable convex compact set

$$\mathcal{Z} \subset \mathcal{Z}^+ = \{Z \in \mathbf{S}^{n+1} : Z \succeq 0, Z_{n+1,n+1} = 1\}$$

the smaller the better, such that $u \in U \rightarrow [u; 1][u; 1]^T \in \mathcal{Z}$

- The ideal selection is

$$\mathcal{Z} = \mathcal{Z}[U] = \text{Conv}\{[u; 1][u; 1]^T : u \in U\}$$

However: $\mathcal{Z}[U]$ usually is computationally intractable.

Important exception:

$$Q \succ 0, U = \{u : u^T Q u \leq 1\} \Rightarrow \mathcal{Z}[U] = \{Z \in \mathcal{Z}^+ : \text{Tr}(ZQ) \leq 1\}$$

“Simple” case: When U is given by S quadratic inequalities:

$$U = \left\{ u \in \mathbb{R}^n : [u; 1]^T Q_s [u; 1] \leq q_s, 1 \leq s \leq S \right\}$$

we can set

$$\mathcal{Z} = \left\{ Z \in \mathbf{S}^{n+1} : Z \succeq 0, Z_{n+1,n+1} = 1, \text{Tr}(Q_s Z) \leq q_s, 1 \leq s \leq S \right\}. (*)$$

Warning: $(*)$ can yield very conservative outer approximation of $\mathcal{Z}[U]$. This conservatism with luck can be reduced by passing from the original description of U to an equivalent one, with emphasis on eliminating/updating linear constraints. For example,

- a constraint of the form $|a^T x - c| \leq r$ should be replaced with
$$(a^T x - c)^2 \leq r^2$$

Note: every linear constraint in the description of U can be written as $\alpha - a^T x \geq 0$ and augmented by redundant constraint $a^T x \geq \beta$, with appropriately selected β . The resulting pair of constraints is equivalent to $|a^T x - c| \leq r$ with $c = \frac{1}{2}[\alpha + \beta]$ and $r = \frac{1}{2}[\alpha - \beta]$.

- It could make sense to write the linear constraints in the description of U in the form $\alpha - a^T x \geq 0$ and add to these constraints their pairwise products.

Quadratic Lifting – Does it Pay?

Situation: Assume that we are given

- convex compact sets $U_\chi \subset \mathbb{R}^{n_\chi}$, $\chi = 1, 2$
- affine mappings $u_\chi \mapsto \mathcal{A}_\chi(u_\chi) : \mathbb{R}^{n_\chi} \rightarrow \mathbb{R}^d$, $\chi = 1, 2$
- convex compact sets $\mathcal{V}_\chi \subset \text{int } \mathbf{S}_+^d$, $\chi = 1, 2$.

These data define families $\mathcal{G}_\chi$ of Gaussian distributions:

$$\mathcal{G}_\chi = \{\mathcal{N}(\mathcal{A}_\chi(u_\chi), \Theta_\chi) : u_\chi \in U_\chi, \Theta_\chi \in \mathcal{V}_\chi\}$$

We can build two types of detectors for $\mathcal{G}_1, \mathcal{G}_2$:

- Affine detector ϕ_{aff} yielded by the solution to the saddle point problem

$$\text{SadVal}_{\text{aff}} = \min_{h \in \mathbb{R}^d} \max_{\substack{u_1 \in U_1, u_2 \in U_2 \\ \Theta_1 \in \mathcal{V}_1, \Theta_2 \in \mathcal{V}_2}} \frac{1}{2} \left[h^T [\mathcal{A}_2(u_2) - \mathcal{A}_1(u_1)] + \frac{1}{2} h^T [\Theta_1 + \Theta_2] h \right],$$

$$\text{Risk}[\phi_{\text{aff}} | \mathcal{G}_1, \mathcal{G}_2] \leq \exp\{\text{SadVal}_{\text{aff}}\}$$

- Quadratic detector ϕ_{lift} yielded by the solution to the saddle point problem

$$\text{SadVal}_{\text{lift}} = \min_{(h, H) \in \mathcal{H}} \max_{\substack{\Theta_1 \in \mathcal{V}_1 \\ \Theta_2 \in \mathcal{V}_2}} \frac{1}{2} \left[\Phi_{\mathcal{A}_1, \mathcal{Z}_1}(-h, -H; \Theta_1) + \Phi_{\mathcal{A}_2, \mathcal{Z}_2}(h, H; \Theta_2) \right],$$

$$\text{Risk}[\phi_{\text{lift}} | \mathcal{G}_1, \mathcal{G}_2] \leq \exp\{\text{SadVal}_{\text{lift}}\}$$

Observation: Assume that sets $\mathcal{V}_\chi$ contain $\succeq$ -largest elements. Then with proper selection of the “design parameters” $\mathcal{Z}_\chi, \Theta_*^{(\chi)}$ participating in the construction of $\Phi_{\mathcal{A}_\chi, \mathcal{Z}_\chi}$, $\chi = 1, 2$, passing from affine to quadratic detectors helps:

$$\text{SadVal}_{\text{lift}} \leq \text{SadVal}_{\text{aff}}$$

Numerical illustration:

- $U_1 = U_1^\rho = \{u \in \mathbb{R}^{12} : u_i \geq \rho, 1 \leq i \leq 12\},$
- $U_2 = U_2^\rho = -U_1^\rho,$
- $A_1 = A_2 \in \mathbb{R}^{8 \times 13};$
- $\mathcal{V}_\chi = \{\Theta_*^{(\chi)} = \sigma_\chi^2 I_8\}$

ρ	σ_1	σ_2	unrestricted H and h	$H = 0$	$h = 0$
0.5	2	2	0.31	0.31	1.00
0.5	1	4	0.24	0.39	0.62
0.01	1	4	0.41	1.00	0.41

Risk of quadratic detector $\phi(\zeta) = h^T \zeta + \frac{1}{2} \zeta^T H \zeta + \varkappa$

- *when deciding on families of Gaussian distributions with common covariance matrix and expectations varying in associated with the families convex sets, passing from affine to quadratic detectors does not help*
- *in general, both affine and purely quadratic components in a quadratic detector are useful.*
- *when deciding on families of Gaussian distributions in the case where distributions from different families can have close expectations, quadratic detectors are useful, while affine ones are not.*

Illustration: Simple Change Point Detection

1

2

3

4

5

6

7

8

15

20

28

34

Frames from a noisy “movie”
When the picture starts to change?

Model: We observe one by one “vectors” (2D images)

$$\omega_t = x_t + \xi_t,$$

- x_t : deterministic image
- $\xi_t \sim \mathcal{N}(0, \sigma^2 I_d)$: independent across observation noises.

Note: We know the range $[\underline{\sigma}, \bar{\sigma}]$ of σ , but perhaps do not know σ exactly.

- We know that $x_1 = x_2$ and want to check whether $x_1 = \dots = x_K$ (“no change”) or there is a change.

Goal: Given an upper bound $\epsilon > 0$ on the probability of false alarm, we want to design a sequential change detection routine capable to detect change, if any.

Approach:

- Pass from observations ω_t , $1 \leq t \leq K$, to observations

$$\zeta_t = \omega_t - \omega_1 = \underbrace{x_t - x_1}_{y_t} + \underbrace{\xi_t - \xi_1}_{\eta_t}, \quad 2 \leq t \leq K$$

- Test hypothesis $H_0 : y_2 = \dots = y_K = 0$ vs. alternative

$$\bigcup_{k=2}^K H_k^\rho, \quad H_k^\rho : y_2 = \dots = y_{k-1} = 0, \quad \|y_k\|_2 \geq \rho$$

via our machinery for testing

red hypothesis H_0

vs.

brown hypotheses $H_2^\rho, \dots, H_K^\rho$

via quadratic liftings $\zeta_t \zeta_t^T$ of observations ζ_t up to closeness

$\mathcal{C}$: all brown hypotheses are close to each other and are not close to the red hypothesis

- We intend to find the smallest ρ for which the $\mathcal{C}$ -risk of the resulting inference is $\leq \epsilon$, and utilize this inference in change point detection.

How It Works

Setup: $\dim y = 256^2 = 65536$, $\bar{\sigma} = 10$, $\bar{\sigma}^2/\underline{\sigma}^2 = 2$, $K = 9$, $\epsilon = 0.01$

Detector: At time $t = 2, \dots, K$, compute

$$\phi_*(\zeta_t) = -2.7138 \frac{\|\zeta_t\|_2^2}{10^5} + 366.9548.$$

- $\phi_*(\zeta_t) < 0 \Rightarrow$ conclude that the change took place and terminate
 $\phi_*(\zeta_t) \geq 0 \Rightarrow$ conclude that there was no change so far and proceed to the next image, if any

- When red hypothesis H_0 holds true, the probability not to claim change on time horizon $2, \dots, K$ is at least 0.99.
- When a brown hypothesis H_k^ρ holds true, the change at time $\leq K$ is detected with probability at least 0.99, provided $\rho \geq \rho_* = 2716.6$ (average per pixel energy in y_k at least by 12% larger than $\bar{\sigma}^2$)
- No test can 0.99-reliably decide via $\zeta_1, \dots, \zeta_k$ on H_k^ρ vs. H_0 , provided $\rho/\rho_* < 0.965$.
- In the movie, the change which takes place at $t = 3$ is detected at $t = 4$.

ESTIMATING SIGNALS IN GAUSSIAN O.S. AND BEYOND

- *Problem of interest*
- *Developing tools, Optimization*
 - *Conic Programming*
 - *Conic Duality*
- *Developing tools, Statistics*
 - *Gauss-Markov Theorem*
- *Optimizing linear estimates*
- *Near-optimality of linear estimates*
- *Byproduct on Semidefinite Relaxation*

Situation: “In the nature” there exists a signal x known to belong to a given convex compact set $\mathcal{X} \subset \mathbb{R}^n$. We observe corrupted by noise affine image of the signal:

$$\omega = Ax + \sigma\xi \in \Omega = \mathbb{R}^m$$

- A : given $m \times n$ sensing matrix
- ξ : random observation noise

Goal: To recover the image Bx of x under a given linear mapping

- B : given $\nu \times n$ matrix.

Risk of a candidate estimate $\hat{x}(\cdot) : \Omega \rightarrow \mathbb{R}^\nu$ is defined as

$$\text{Risk}[\hat{x}|\mathcal{X}] = \sup_{x \in \mathcal{X}} \sqrt{\mathbf{E}_\xi \left\{ \|Bx - \hat{x}(Ax + \sigma\xi)\|_2^2 \right\}}$$

Risk^2 is the worst-case, over $x \in \mathcal{X}$, expected $\|\cdot\|_2^2$ recovery error.

Ideal goal: To build the minmax optimal estimate $\hat{x}_*$ achieving the *minmax risk*

$$\text{RiskOpt}[\mathcal{X}] = \inf_{\hat{x}(\cdot)} \text{Risk}[\hat{x}|\mathcal{X}].$$

Agenda: Under appropriate assumptions on $\mathcal{X}$, we show that

- *One can build, in a computationally efficient fashion, (nearly) the best, in terms of risk, estimate in the family of linear estimates*

$$\hat{x}(\omega) = \hat{x}_H(\omega) = H^T \omega \quad [H \in \mathbb{R}^{m \times \nu}]$$

- *The resulting linear estimate is nearly optimal among all estimates, linear and nonlinear alike.*

Why linear estimates?

As it was announced, a “nearly optimal” linear estimate can be built in a computationally efficient fashion.

In contrast, Exactly minimax optimal estimate is *unknown* even in the simplest case when the observation is

$$\omega = x + \eta$$

with $\eta \sim \mathcal{N}(0, \sigma^2)$ and $x \in \mathcal{X} = [-1, 1]$

Note: The statistical “magic wand” – Maximum Likelihood Estimate – can be *heavily nonoptimal* already in the Simple case.

blue: $\mathcal{X}$ magenta: $A\mathcal{X}$

- $\mathcal{X} = \{x \in \mathbb{R}^n : x_n^2 + \epsilon^{-2} \sum_{i=1}^{n-1} x_i^2 \leq 1\}$
 - $A = \text{Diag}\{1/\epsilon, \dots, 1/\epsilon, 1\}, \quad \eta \sim \mathcal{N}(0, \sigma^2 I_n), \quad B = I_n$
- $\Rightarrow$ MLE: $\hat{x}_{ml}(\omega) = A^{-1} \cdot \text{argmin}_{\|u\|_2 \leq 1} \|\omega - u\|_2$

When $\sigma \ll 1$, $\sigma^2 n \geq O(1)$, and $\epsilon \leq O(\sigma)$, the risk of MLE is $O(1)$, while the risk of the linear estimate $\hat{x}(\omega) = \omega_n$ is $O(\sigma) \ll O(1)$.

$$\omega = Ax + \sigma\xi \text{ \& } x \in \mathcal{X}' \Rightarrow \hat{x}_H(\omega) := H^T\omega \approx Bx$$

Design of linear estimates: challenges

Fact: Unless A and B are simple – just diagonal, – analytical risk analysis/design of linear estimates is problematic even in the **Simple case** where $\xi \sim \mathcal{N}(0, I)$ and $\mathcal{X} = \{x \in \mathbb{R}^n : \sum_i a_i^2 x_i^2 \leq 1\}$ is an ellipsoid.

Reason: *Optimal minimax risk is determined by difficult to represent and analyze interplay between “geometries” of $A, B, \mathcal{X}$:*

- on one hand, poor conditioning of A can make impossible good recovery of some components of x even in a low noise
- on the other hand, the geometries of $\mathcal{X}$ and/or B can make “difficult to recover” components of x irrelevant – these components can be a priori small due to the geometry of the signal set $\mathcal{X}$, or can be suppressed by B due to the geometry of B .

Surprisingly, difficulties basically disappear when passing from *descriptive analytical* to *computationally efficient operational* design and risk analysis of estimates. Specifically, in *Simple case with no assumptions on A, B* (and even far beyond) one can build in a *computationally efficient* fashion *provably minimax optimal*, up to logarithmic factors, linear estimates.

Developing tools, Optimization: “structure-revealing” representation of Convex Problems: Conic Programming

- When passing from a Linear Programming program

$$\min_x \{c^T x : Ax - b \geq 0\} \quad (*)$$

to a nonlinear convex one, the traditional wisdom is to replace linear inequality constraints

$$a_i^T x - b_i \geq 0$$

with nonlinear ones:

$$g_i(x) \geq 0 \quad [g_i \text{ are concave}]$$

- There exists, however, another way to introduce nonlinearity, namely, to replace the coordinate-wise *vector* inequality

$$y \geq z \Leftrightarrow y - z \in \mathbb{R}_+^m = \{u \in \mathbb{R}^m : u_i \geq 0 \forall i\} \quad [y, z \in \mathbb{R}^m]$$

with another *vector* inequality

$$y \geq_{\mathbf{K}} z \Leftrightarrow y - z \in \mathbf{K} \quad [y, z \in \mathbb{R}^m]$$

where $\mathbf{K}$ is a *regular cone* (i.e., closed, pointed and convex cone with a nonempty interior) in $\mathbb{R}^m$.

$$y \geq_{\mathbf{K}} z \Leftrightarrow y - z \in \mathbf{K} \quad [y, z \in \mathbb{R}^m]$$

$\mathbf{K}$: closed, pointed and convex cone in $\mathbb{R}^m$ with a nonempty interior.

Requirements on $\mathbf{K}$ ensure that $\geq_{\mathbf{K}}$ obeys the usual rules for inequalities:

- $\geq_{\mathbf{K}}$ is a *partial order*:

$$x \geq_{\mathbf{K}} x \quad \forall x \quad \text{[reflexivity]}$$

$$(x \geq_{\mathbf{K}} y \ \& \ y \geq_{\mathbf{K}} x) \Rightarrow x = y \quad \text{[antisymmetry]}$$

$$(x \geq_{\mathbf{K}} y, \ y \geq_{\mathbf{K}} z) \Rightarrow x \geq_{\mathbf{K}} z \quad \text{[transitivity]}$$

- $\geq_{\mathbf{K}}$ is compatible with linear operations: the validity of $\geq_{\mathbf{K}}$ inequality is preserved when we multiply both sides by the same nonnegative real and add to it another valid $\geq_{\mathbf{K}}$ -inequality;
- in a sequence of $\geq_{\mathbf{K}}$ -inequalities, one can pass to limits:

$$\{a_i \geq_{\mathbf{K}} b_i, \ i = 1, 2, \dots \ \& \ a_i \rightarrow a \ \& \ b_i \rightarrow b\} \Rightarrow a \geq_{\mathbf{K}} b$$

$$y \geq_{\mathbf{K}} z \Leftrightarrow y - z \in \mathbf{K} \quad [y, z \in \mathbb{R}^m]$$

$\mathbf{K}$: closed, pointed and convex cone in $\mathbb{R}^m$ with a nonempty interior.

- one can define the strict version $>_{\mathbf{K}}$ of $\geq_{\mathbf{K}}$:

$$a >_{\mathbf{K}} b \Leftrightarrow a - b \in \text{int } \mathbf{K}.$$

Arithmetics of $>_{\mathbf{K}}$ and $\geq_{\mathbf{K}}$ inequalities is completely similar to the arithmetics of the usual coordinate-wise $>$ and $\geq$.

LP problem:

$$\min_x \{c^T x : Ax - b \geq 0\} \Leftrightarrow \min_x \{c^T x : Ax - b \in \mathbb{R}_+^m\}$$

General Conic problem:

$$\min_x \{c^T x : Ax - b \succeq_{\mathbf{K}} 0\} \Leftrightarrow \min_x \{c^T x : Ax - b \in \mathbf{K}\}$$

- (A, b) – *data* of conic problem
- $\mathbf{K}$ - structure of conic problem

Remarks:

- Every convex problem admits equivalent conic reformulation
- With conic formulation, convexity is “built in”; with the standard MP formulation convexity should be kept in mind as an additional property.

(??) A general convex cone has no more structure than a general convex function. Why conic reformulation is “structure-revealing”?

(!!) Just 3 types of cones allow to represent an extremely wide spectrum (“essentially all”) of convex problems!

$$\min_x \{c^T x : Ax - b \geq_{\mathbf{K}} 0\} \Leftrightarrow \min_x \{c^T x : Ax - b \in \mathbf{K}\}$$

Three Magic Families of cones:

- *LP: Nonnegative orthants* $\mathbb{R}_+^m$ – direct products of m nonnegative rays
 $\mathbb{R}_+ = \{s \in \mathbb{R} : s \geq 0\}$

giving rise to Linear Programming programs

$$\min_s \{c^T x : a_\ell^T x - b_\ell \geq 0, 1 \leq \ell \leq q\}.$$

- *CQP: Direct products of Lorentz cones*

$$\mathbf{L}_+^p = \{u \in \mathbb{R}^p : u_p \geq (\sum_{i=1}^{p-1} u_i^2)^{1/2}\}$$

giving rise to Conic Quadratic programs

$$\min_x \{c^T x : \|A_\ell x - b_\ell\|_2 \leq c_\ell^T x - d_\ell, 1 \leq \ell \leq q\}.$$

- *SDP: Direct products of Semidefinite cones*

$$\mathbf{S}_+^p = \{M \in \mathbf{S}^p : M \succeq 0\}$$

giving rise to Semidefinite programs

$$\min_x \left\{ c^T x : \underbrace{\lambda_{\min}(\mathcal{A}^\ell(x))}_{\Leftrightarrow \mathcal{A}^\ell(x) \succeq 0} \geq 0, 1 \leq \ell \leq q \right\}.$$

where $\mathbf{S}^p$ is the space of $p \times p$ real symmetric matrices, $\mathcal{A}_\ell(x) \in \mathbf{S}^p$ are affine in x and $\lambda_{\min}(S)$ is the minimal eigenvalue of $S \in \mathbf{S}^p$.

Note: Constraint $\mathcal{A}(x) \succeq 0$ stating that a symmetric matrix affinely depending on decision variables is $\succeq 0$ is called LMI – Linear Matrix Inequality.

For details, see https://www2.isye.gatech.edu/~nemirovs/lmco_run.pdf

What can be reduced to $\mathcal{LP}/\mathcal{CQP}/\mathcal{SDP}$? Calculus of Conic Programs

Let $\mathcal{K}$ be a family of regular cones closed w.r.t. taking direct products.

Definition: A $\mathcal{K}$ -representation of a set $X \subset \mathbb{R}^n$ is a representation

$$X = \{x \in \mathbb{R}^n : \exists u \in \mathbb{R}^m : Ax + Bu - b \in \mathbf{K}\} \quad (*)$$

where $\mathbf{K} \in \mathcal{K}$. X is called $\mathcal{K}$ -representable, if X admits a $\mathcal{K}$ -r.

Note: Minimizing a linear objective $c^T x$ over a $\mathcal{K}$ -representable set X reduces to a conic program on a cone from $\mathcal{K}$.

Indeed, given (*), problem $\min_{x \in X} c^T x$ is equivalent to

$$\text{Opt} = \min_{x,u} \{c^T x : Ax + Bu - b \in \mathbf{K}\}$$

Definition: $\mathcal{K}$ -representation of a function $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is a $\mathcal{K}$ -representation of the epigraph of f :

$\text{Epi}\{f\} := \{(x, t) : t \geq f(x)\} = \{x, t : \exists v : Px + pt + Qv - q \in \mathbf{K}\}, \quad \mathbf{K} \in \mathcal{K}$
 f is called $\mathcal{K}$ -representable, if f admits a $\mathcal{K}$ -r.

- A level set of a $\mathcal{K}$ -r. function is $\mathcal{K}$ -r.:

$$\begin{aligned} \text{Epi}\{f\} &:= \{(x, t) : t \geq f(x)\} = \{x, t : \exists v : Px + pt + Qu - q \in \mathbf{K}\} \\ &\Rightarrow \{x : f(x) \leq c\} = \{x : \exists v : Px + Qu - [q - cp] \in \mathbf{K}\} \end{aligned}$$

- Minimization of a $\mathcal{K}$ -r. function f over a $\mathcal{K}$ -r. set X reduces to a conic program on a cone from $\mathcal{K}$:

$$\begin{aligned} &\left\{ \begin{array}{ll} x \in X & \Leftrightarrow \exists u : Ax + Bu - b \in \mathbf{K}_X \\ t \geq f(x) & \Leftrightarrow \exists v : Px + pt + Qv - q \in \mathbf{K}_f \end{array} \right\} \\ &\quad \Downarrow \\ &\quad \min_{x \in X} f(x) \\ &\quad \Updownarrow \\ &\min_{t, x, u, v} \left\{ t : [Ax + Bu - b; Px + pt + Qv - q] \in \underbrace{\mathbf{K}_X \times \mathbf{K}_f}_{\in \mathcal{K}} \right\} \end{aligned}$$

- Investigating “expressive abilities” of generic Magic conic problems reduces to answering the question

What are $\mathcal{LP}/\mathcal{CQP}/\mathcal{SDP}$ -r. functions/sets?

- *“Built-in” restriction is convexity: A $\mathcal{K}$ -representable set/function must be convex.*

Good news: *Convexity is, essentially, **the only** restriction: for all practical purposes, all convex sets/functions arising in applications are $\mathcal{SDP}$ -r. Quite rich families of convex functions/sets are $\mathcal{LP}/\mathcal{CQP}$ -r.*

Note: Nonnegative orthants are direct products of (1-dimensional) Lorentz cones, and Lorentz cones are intersections of semidefinite cones and properly selected linear subspaces $\Rightarrow \mathcal{LP} \subset \mathcal{CQP} \subset \mathcal{SDP}$.

Let $\mathcal{K}$ be a family of regular cones closed w.r.t. taking direct products and passing from a cone $\mathbf{K}$ to its *dual cone*:

$$\mathbf{K}_* = \{\lambda : \langle \lambda, \xi \rangle \geq 0 \ \forall \xi \in \mathbf{K}\}$$

Note: $\mathbf{K}_*$ is regular cone provided $\mathbf{K}$ is so, and

$$(\mathbf{K}_*)_* = \mathbf{K}$$

Fact: $\mathcal{K}$ -representable sets/functions admit fully algorithmic calculus:

- all basic convexity-preserving operations with functions/sets, as applied to $\mathcal{K}$ -r. operands, produce $\mathcal{K}$ -r. results, and the resulting $\mathcal{K}$ -r.'s are readily given by $\mathcal{K}$ -r.'s of the operands.
- “Calculus rules” are independent of what $\mathcal{K}$ is.

$\Rightarrow$ Starting with “raw materials” (characteristic for $\mathcal{K}$ elementary $\mathcal{K}$ -r. sets/functions) and applying calculus rules, we can recognize $\mathcal{K}$ -representability and get explicit $\mathcal{K}$ -r.'s of sets/functions of interest.

“Calculus of $\mathcal{K}$ -representability” basics:

Sets: If $X_1, \dots, X_k$ are $\mathcal{K}$ -r. sets, so are their

- *intersections,*
- *images under affine mappings,*
- *direct products,*
- *inverse images under affine mappings.*

Functions: If $f_1, \dots, f_k$ are $\mathcal{K}$ -r. functions, so are their

- *linear combinations with nonnegative coefficients,*
- *superpositions with affine mappings.*

Moreover, *if $F, f_1, \dots, f_k$ are $\mathcal{K}$ -r. functions, so is the superposition $F(f_1(x), \dots, f_k(x))$ provided that F is monotonically nondecreasing in its arguments.*

More advanced convexity-preserving operations preserve $\mathcal{K}$ -representability under (pretty mild!) regularity conditions. This includes

For sets: taking *conic hulls* and *convex hulls of (finite) unions* and passing from a set to its *recessive cone*, or *polar*, or *support function*

For functions: *partial minimization*, *projective transformation*, and *taking Fenchel dual*.

Note: *Calculus rules are simple and algorithmic $\Rightarrow$ Calculus can be run in a compiler [used in CVX].*

Illustration

$\min c^T x + d^T y$ $y \geq 0, Ax + By \leq b$	
$2y_1^{-\frac{7}{2}}y_2^{-3}y_3^{-\frac{1}{5}} + 3y_2^{-\frac{3}{2}}y_4^{-\frac{2}{3}} \leq e^T x + 4y_1^{\frac{1}{5}}y_2^{\frac{2}{5}}y_3^{\frac{2}{5}} + 5y_3^{\frac{1}{3}}y_4^{\frac{2}{5}}$	$\begin{bmatrix} x_1 - x_2 & x_3 + x_2 & & & \\ x_3 + x_2 & x_2 - x_4 & x_5 - 6 & & \\ & x_5 - 6 & x_6 + x_7 & -x_8 & \\ & & -x_8 & x_5 & \end{bmatrix} \succeq 0$
$\text{Det} \left(\begin{bmatrix} x_1 & x_2 & x_3 & x_4 & x_5 \\ x_2 & x_6 & x_7 & x_8 & x_9 \\ x_3 & x_7 & x_{10} & x_{11} & x_{12} \\ x_4 & x_8 & x_{11} & x_{13} & x_{14} \\ x_5 & x_9 & x_{12} & x_{14} & x_{15} \end{bmatrix} \right) \geq 1$	
$\text{Sum of 2 largest singular values of } \begin{bmatrix} x_1 & x_2 & x_3 \\ x_4 & x_5 & x_6 \\ x_7 & x_8 & x_9 \\ x_{10} & x_{11} & x_{12} \\ x_{13} & x_{14} & x_{15} \end{bmatrix} \text{ is } \leq 6$	
$1 - \sum_{i=1}^6 [x_i - x_{i+1}] s^i \leq 0, \quad \frac{3}{2} \leq s \leq 6$ $\sum_{i=1}^4 x_{2i} \cos(i\phi) - \sum_{i=1}^4 x_i \sin(i\phi) \leq 1, \quad \frac{\pi}{3} \leq \phi \leq \frac{\pi}{2}$	

- blue part is in $\mathcal{LP}$
- blue-magenta part of the problem is in $\mathcal{CQP}$ (can be approximated, *in a polynomial time*, by $\mathcal{LP}$)
- entire problem is in $\mathcal{SDP}$
- reductions to $\mathcal{LP}/\mathcal{CQP}/\mathcal{SDP}$ are “fully algorithmic.”

Conic Duality

Conic Program admits nice *Duality Theory* completely similar to LP Duality.

Primal problem:

$$\min_x \{c^T x : Ax - b \geq_{\mathbf{K}} 0, Rx = r\}$$

$\Updownarrow$ [passing to primal slack $\xi = Ax - b$]

$$\min_{\xi} \{e^T \xi : \xi \in [\mathcal{L} - b] \cap \mathbf{K}\}$$

$$[e : A^T e + R^T f = c \text{ for some } f, \quad \mathcal{L} = \{Au : Ru = 0\}]$$

Dual problem:

$$\max_{y,z} \{b^T y : A^T y + R^T z = c, y \geq_{\mathbf{K}_*} 0\}$$

$\Updownarrow$

$$\max_y \{b^T y : y \in [\mathcal{L}^\perp + e] \cap \mathbf{K}_*\}$$

$[\mathbf{K}_* : \text{cone dual to } \mathbf{K}]$

Primal problem:

$$\min_{\xi} \left\{ e^T \xi : \xi \in [\mathcal{L} - b] \cap \mathbf{K} \right\}$$
$$\left[e : A^T e + R^T f = c \text{ for some } f, \quad \mathcal{L} = \{ Au : Ru = 0 \} \right]$$

Dual problem:

$$\max_y \left\{ b^T y : y \in [\mathcal{L}^\perp + e] \cap \mathbf{K}_* \right\}$$
$$[\mathbf{K}_* : \text{cone dual to } \mathbf{K}]$$

Note:

- the dual problem is conic along with primal
- the duality is completely symmetric
- Cones from Magic Families are self-dual, so that the dual of a Linear/Conic Quadratic/Semidefinite program is of exactly the same type.

Deriving the dual problem

Primal problem:

$$\text{Opt}(P) = \min_x \left\{ c^T x : \begin{array}{l} A_i x - b_i \in \mathbf{K}^i, \ i \leq m \\ Rx = r \end{array} \right\} \quad (P)$$

Goal: find a systematic way to bound $\text{Opt}(P)$ *from below*.

Observation: When $y_i \in \mathbf{K}_*^i$, the scalar inequality $y_i^T A_i x \geq y_i^T b_i$ is a consequence of the constraint $A_i x - b_i \in \mathbf{K}^i$.

If z is a vector of the same dimension as r , the scalar inequality $z^T Rx \geq z^T r$ is a consequence of the constraint $Rx = r$.

$\Rightarrow$ Whenever $y_i \in \mathbf{K}_*^i$ for all i and z is a vector of the same dimension as r , the scalar linear inequality

$$[\sum_i A_i^T y_i + R^T z]^T x \geq \sum_i b_i^T y_i + r^T z$$

is a consequence of the constraints in (P)

$\Rightarrow$ Whenever $y_i \in \mathbf{K}_*^i$ for all i and z is a vector of the same dimension as r such that

$$\sum_i A_i^T y_i + R^T z = c,$$

the quantity $\sum_i b_i^T y_i + r^T z$ is a lower bound on $\text{Opt}(P)$.

Dual problem:

$$\text{Opt}(D) = \max_{y_i, z} \left\{ \sum_i b_i^T y_i + r^T z : \begin{array}{l} y_i \in \mathbf{K}_*^i, i \leq m \\ \sum_i A_i^T y_i + R^T z = c \end{array} \right\} \quad (D)$$

is exactly the problem of maximizing this lower bound on $\text{Opt}(P)$.

Definition A conic problem

$$\min_x \left\{ c^T x : \begin{array}{l} A_i x - b_i \in \mathbf{K}^i, \ i \leq m \\ Ax \geq b \\ Rx = r \end{array} \right\} \quad (C)$$

is called *strictly feasible*, if there exists a *feasible* solution $\bar{x}$ where *all conic and $\leq$ constraints* are *strictly satisfied*:

$$A_i \bar{x} - b_i \in \text{int } \mathbf{K}^i \ \forall i \ \& \ A \bar{x} > b.$$

It is called *essentially strictly feasible*, if there exists a *feasible* solution $\bar{x}$ where *all non-polyhedral constraints* are strictly satisfied:

$$A_i \bar{x} - b_i \in \text{int } \mathbf{K}^i \ \forall i.$$

Conic Duality Theorem Consider a conic problem

$$\text{Opt}(P) = \min_x \left\{ c^T x : \begin{array}{l} A_i x - b_i \in \mathbf{K}^i, \ i \leq m \\ R x = r \end{array} \right\} \quad (P)$$

along with its dual

$$\text{Opt}(D) = \max_{y_i, z} \left\{ \sum_i b_i^T y_i + r^T z : \begin{array}{l} y_i \in \mathbf{K}_*^i, \ i \leq m \\ \sum_i A_i^T y_i + R^T z = c \end{array} \right\} \quad (D)$$

Then

- [Symmetry] Duality is symmetric: the dual problem is conic, and its dual is (equivalent to) the primal problem;
- [Weak duality] One has $\text{Opt}(D) \leq \text{Opt}(P)$;
- [Strong duality] Let one of the problems be *essentially strictly feasible* and bounded. Then the other problem is solvable, and

$$\text{Opt}(D) = \text{Opt}(P).$$

In particular, if both problems are *essentially strictly feasible*, both are solvable with equal optimal values.

$$\begin{aligned} \min_x \left\{ c^T x : \begin{array}{l} A_i x - b_i \in \mathbf{K}^i, i \leq m \\ R x = r \end{array} \right\} & \quad (P) \\ \updownarrow & \\ \max_{y_i, z} \left\{ \sum_i b_i^T y_i + r^T z : \begin{array}{l} y_i \in \mathbf{K}_*^i, i \leq m \\ \sum_i A_i^T y_i + R^T z = c \end{array} \right\} & \quad (D) \end{aligned}$$

Conic Programming Optimality Conditions

Let both (P) and (D) be essentially strictly feasible. Then a pair $(x, [\{y_i\}, z])$ of primal and dual *feasible* solutions is comprised of optimal solutions to the respective problems if and only if

[Zero Duality Gap] $\text{DualityGap}(x, [\{y_i\}, z]) := c^T x - [\sum_i b_i^T y_i + r^T z] = 0$

Indeed,

$$\text{DualityGap}(x, [\{y_i\}, z]) = \underbrace{[c^T x - \text{Opt}(P)]}_{\geq 0} + \underbrace{[\text{Opt}(D) - [\sum_i b_i^T y_i + r^T z]]}_{\geq 0}$$

[Complementary Slackness] $[A_i x - b_i]^T y_i = 0, i \leq m$

Indeed,

$$\begin{aligned} \sum_i \underbrace{[A_i x - b_i]^T y_i}_{\geq 0} &= [\sum_i A_i^T y_i] x - \sum_i b_i^T y_i = \underbrace{[c - R^T z]^T x}_{=c^T x - r^T z} - \sum_i b_i^T y_i = c^T x - [\sum_i b_i^T y_i + r^T z] \\ &= \text{DualityGap}(x, [\{y_i\}, z]) \end{aligned}$$

- Conic Duality, same as the LP one, is
 - *fully algorithmic*: to write down the dual, given the primal, is a purely mechanical process
 - *fully symmetric*: the dual problem “remembers” the primal one
- Cf. *Lagrange Duality*:

$$\min_x \{f(x) : g_i(x) \leq 0, i = 1, \dots, m\} \quad (P)$$

$\Downarrow$

$$\max_{y \geq 0} \underline{L}(y) \quad (D)$$

$$\left[\underline{L}(y) = \min_x \left\{ f(x) + \sum_i y_i g_i(x) \right\} \right]$$

- Dual “exists in the nature”, but is given implicitly; its objective, typically, is not available in a closed form
- Duality is asymmetric: given $\underline{L}(\cdot)$, we, typically, cannot recover f and $g_i \dots$

Developing Tools, Optimization: Schur Complement Lemma

Lemma Symmetric block matrix $\left[\begin{array}{c|c} P & S^T \\ \hline S & R \end{array} \right]$ with $R \succ 0$ is positive semidefinite if and only if the matrix $P - S^T R^{-1} S$ is so.

Proof: $\left[\begin{array}{c|c} P & S^T \\ \hline S & R \end{array} \right] \succeq 0$ iff

$$\begin{aligned} 0 &\leq \min_{u,v} [u^T P u + 2u^T S^T v + v^T R v] \\ &= \min_u \left[\underbrace{\min_v [u^T P u + 2u^T S^T v + v^T R v]}_{\text{achieved when } v = -R^{-1} S u} \right] \\ &= \min_u u^T [P - S^T R^{-1} S] u. \end{aligned}$$

□

Developing Tools, Statistics: Bayesian Risk and Gauss-Markov Theorem

An alternative to the *worst-case risk* (the *worst* over the signals of interest, performance of a statistical procedure) is *average performance*, with the average taken over some *prior* probability distribution on the signals.

In the problem of $\|\cdot\|_2$ -recovering Bx via noisy observation

$$\omega = Ax + \sigma\xi, \quad \xi \sim P$$

this alternative reads as follows:

(!) *Given a probability distribution Π of signal $x \in \mathbb{R}^n$, find an estimate $\hat{x}(\cdot)$ which minimizes*

$$\text{Risk}^2[\hat{x}|\Pi] := \int_{\Pi} \left\{ \int_{\mathbb{R}^m} \|Bx - \hat{x}(Ax + \sigma\xi)\|_2^2 P(d\xi) \right\} \Pi(dx)$$

– *the average, over the distribution Π of signals x , of expected $\|\cdot\|_2^2$ recovery error of Bx via observation $Ax + \sigma\xi$.*

(!) Given a probability distribution Π of signal $x \in \mathbb{R}^n$, find an estimate $\hat{x}(\cdot)$ which minimizes

$$\text{Risk}^2[\hat{x}|\Pi] := \int_{\Pi} \left\{ \int_{\mathbb{R}^m} \|Bx - \hat{x}(Ax + \sigma\xi)\|_2^2 P(d\xi) \right\} \Pi(dx)$$

– the average, over the distribution Π of signals x , of expected $\|\cdot\|_2^2$ recovery error of Bx via observation $Ax + \sigma\xi$.

- Q : joint distribution of $(x, \omega = Ax + \sigma\xi)$ induced by distributions Π and P of independent r.v.'s x, ξ
- R_ω : conditional distribution of x , given ω of x ,
- W : marginal distribution of ω .

We have

$$\begin{aligned} \text{Risk}^2[\hat{x}|\Pi] &= \int_{\mathbb{R}_x^n \times \mathbb{R}_\omega^m} \|Bx - \hat{x}(\omega)\|_2^2 Q(dx, d\omega) \\ &= \int_{\mathbb{R}^m} \left\{ \int_{\mathbb{R}^n} \|Bx - \hat{x}(\omega)\|_2^2 R_\omega(dx) \right\} W(d\omega) \end{aligned}$$

Evident Fact: Assuming that a probability distribution S on $\mathbb{R}^n$ possesses finite second moments, one has

$$\min_{c \in \mathbb{R}^p} \int_{\mathbb{R}^n} \|Bx - c\|_2^2 S(dx) = \int_{\mathbb{R}^n} \|Bx - c_*\|_2^2 S(dx),$$

$$c_* = \int_{\mathbb{R}^n} Bx S(dx).$$

$\Rightarrow$ An optimal, in terms of $\text{Risk}^2[\hat{x}|\Pi]$, estimate $\hat{x}(\cdot)$ of Bx via $\omega = Ax + \sigma\xi$ is

$$\hat{x}_*(\omega) = \int_{\mathbb{R}^n} Bx R_\omega(dx)$$

Gauss-Markov Theorem [Normal Correlation Theorem] Let $x \in \mathbb{R}^n$ and $\xi \in \mathbb{R}^m$ be independent zero mean Gaussian random vectors. Assuming $\sigma > 0$ and the covariance matrix of ξ to be positive definite, an optimal solution $\hat{x}(\cdot)$ to the risk minimization problem

$$\min_{\hat{x}(\cdot)} \mathbf{E}_{x,\xi} \left\{ \|Bx - \hat{x}(Ax + \sigma\xi)\|_2^2 \right\}$$

exists and can be selected as a linear function of $\omega = Ax + \sigma\xi$.

Under the premise of Gauss-Markov Theorem, with $x \sim \mathcal{N}(0, Q)$ and $\xi \sim \mathcal{N}(0, I_m)$, direct computation shows that

$$\begin{aligned}\hat{x}_*(\omega) &= \left[[\sigma^2 I_m + AQA^T]^{-1} AQB^T \right] \omega \\ \text{Risk}^2[\hat{x}_* | \mathcal{N}(0, Q)] &= \text{Tr}(BQB^T - BQA^T [\sigma^2 I_m + AQA^T]^{-1} AQB^T)\end{aligned}$$

Optimizing linear estimates

Situation: “In the nature” there exists a signal x known to belong to a given convex compact set $\mathcal{X} \subset \mathbb{R}^n$. We observe corrupted by noise affine image of the signal:

$$\omega = Ax + \sigma\xi \in \Omega = \mathbb{R}^m$$

where

- A is a given $m \times n$ sensing matrix
- ξ is a random noise

Goal: To recover the image Bx of x where

- B : given $\nu \times n$ matrix.

Risk of a candidate estimate $\hat{x}(\cdot) : \Omega \rightarrow \mathbb{R}^\nu$ is

$$\text{Risk}[\hat{x}|\mathcal{X}] = \sup_{x \in \mathcal{X}} \sqrt{\mathbf{E}_\xi \left\{ \|Bx - \hat{x}(Ax + \sigma\xi)\|_2^2 \right\}}$$

Assumption on noise: ξ is zero mean with unit covariance matrix.

$\Rightarrow$ The risk of a linear estimate $\hat{x}_H(\omega) = H^T \omega$ is given by

$$\begin{aligned}\text{Risk}^2[\hat{x}_H|\mathcal{X}] &= \max_{x \in \mathcal{X}} \mathbf{E}_{\xi} \left\{ \|[B - H^T A]x - \sigma H^T \xi\|_2^2 \right\} \\ &= \max_{x \in \mathcal{X}} \left\{ \|[B - H^T A]x\|_2^2 + \sigma^2 \mathbf{E}_{\xi} \{ \text{Tr}(H^T \xi \xi^T H) \} \right\} \\ &= \sigma^2 \text{Tr}(H^T H) + \underbrace{\max_{x \in \mathcal{X}} x^T [B - H^T A]^T [B - H^T A] x}_{\Psi(H)}.\end{aligned}$$

Note: Ψ is convex $\Rightarrow$ building the minimum risk linear estimate reduces to solving *convex* minimization problem

$$\text{Opt} = \min_H \left[\Psi(H) + \sigma^2 \text{Tr}(H^T H) \right]. \quad (*)$$

However: Convex function Ψ is given implicitly and can be difficult to compute, making $(*)$ difficult as well.

$$\begin{aligned} \text{Opt} &= \min_H \left[\sigma^2 \text{Tr}(H^T H) + \Psi(H) \right] \\ \Psi(H) &= \max_{x \in \mathcal{X}} x^T [B - H^T]^T [B - H^T A] x \end{aligned} \quad (*)$$

Fact: Basically, the only cases when $(*)$ is known to be easy are those when

- $\mathcal{X}$ is given as a convex hull of finite set of moderate cardinality
- $\mathcal{X}$ is an ellipsoid.

When $\mathcal{X}$ is a box, computing Ψ is NP-hard...

- When Ψ is difficult to compute, we can to replace Ψ in the design problem $(*)$ with an efficiently computable convex upper bound $\Psi^+(H)$.

We are about to consider a family of sets $\mathcal{X}$ – *ellitopes* – for which reasonably tight bounds Ψ^+ of desired type are available.

A basic ellitope is a set $\mathcal{Y} \subset \mathbb{R}^N$ given as

$$\mathcal{Y} = \{y \in \mathbb{R}^N : \exists t \in \mathcal{T} : y^T S_k y \leq t_k, k \leq K\}$$

where

- $S_k \succeq 0$ are positive semidefinite matrices with $\sum_k S_k \succ 0$
- $\mathcal{T}$ is a convex compact subset of K -dimensional nonnegative orthant $\mathbb{R}_+^K$ such that
 - $\mathcal{T}$ contains some positive vectors
 - $\mathcal{T}$ is *monotone*: if $0 \leq t' \leq t$ and $t \in \mathcal{T}$, then $t' \in \mathcal{T}$ as well.

An ellitope $\mathcal{X}$ is linear image of a basic ellitope:

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists y \in \mathbb{R}^N, t \in \mathcal{T} : x = Fy, y^T S_k y \leq t_k, k \leq K\}$$

- F is a given $n \times N$ matrix,

Note: *Every ellitope is a symmetric w.r.t. the origin convex compact set.*

Examples of basic ellitopes

A. Ellipsoid centered at the origin

$$(K = 1, \mathcal{T} = [0; 1])$$

B. (Bounded) intersection of K ellipsoids/elliptic cylinders centered at the origin

$$(\mathcal{T} = \{t \in \mathbb{R}^K : 0 \leq t_k \leq 1, k \leq K\})$$

C. Box $\{x \in \mathbb{R}^n : -1 \leq x_i \leq 1\}$

$$(\mathcal{T} = \{t \in \mathbb{R}^n : 0 \leq t_k \leq 1, k \leq K = n\}, x^T S_k x = x_k^2)$$

C'. Dual ball of Vasershtein norm $\mathcal{B}_* = \{u \in \mathbb{R}^n : |u_i - u_j| \leq d_{ij}, i, j \leq n\}$

$$\mathcal{B}_* = \{u \in \mathbb{R}^n : u^T S_{ij} u \leq 1, 1 \leq i \leq j \leq n\}, S_{ij} = \begin{cases} d_{ij}^{-2} [e_i - e_j][e_i - e_j]^T, & i < j \leq n, \\ 4[\max_{i,j} d_{ij}]^{-2} e_i e_i^T, & i = j \leq n \end{cases}$$

where e_i, e_j are canonic orths of $\mathbb{R}^n$.

D. ℓ_p -ball $\mathcal{X} = \{x \in \mathbb{R}^n : \|x\|_p \leq 1\}$ with $p \geq 2$

$$(\mathcal{T} = \{t \in \mathbb{R}_+^n : \|t\|_{p/2} \leq 1\}, x^T S_k x = x_k^2, k \leq K = n)$$

Ellitopes admit fully algorithmic calculus:

if $\mathcal{X}_i$, $1 \leq i \leq I$, are ellitopes, so are their

- intersection $\bigcap_i \mathcal{X}_i$
- direct product $\mathcal{X}_1 \times \dots \times \mathcal{X}_I$
- arithmetic sum $\mathcal{X}_1 + \dots + \mathcal{X}_I$
- linear images $\{Ax : x \in \mathcal{X}_i\}$
- inverse linear images $\{y : Ay \in \mathcal{X}_i\}$ under linear embeddings A

Observation: Let

$$\mathcal{X} = \{x : \exists(t \in \mathcal{T}, y) : x = Fy, y^T S_k y \leq t_k, k \leq K\} \quad (*)$$

be an ellitope. Given a quadratic form $x^T W x$, $W \in \mathbf{S}^n$, we have

$$\max_{x \in \mathcal{X}} x^T W x \leq \min_{\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) : \lambda \geq 0, \sum_{k=1}^K \lambda_k S_k \succeq F^T W F \right\}$$

where

$$\phi_{\mathcal{T}}(\lambda) = \max_{t \in \mathcal{T}} t^T \lambda$$

is the support function of $\mathcal{T}$.

Indeed, assume that $\lambda \geq 0$ and $F^T W F \preceq \sum_k \lambda_k S_k$. By definition of the ellitope, if $x \in \mathcal{X}$,

$$\exists(t \in \mathcal{T}, y) : y^T S_k y \leq t_k \quad \forall k \leq K, \quad x = Fy$$

We have $x^T W x = y^T F^T W F y$, but

$$y^T F^T W F y \leq \sum_k \lambda_k y^T S_k y \leq \sum_k \lambda_k t_k \leq \max_{t \in \mathcal{T}} t^T \lambda = \phi_{\mathcal{T}}(\lambda),$$

and we conclude that

$$x^T W x \leq \phi_{\mathcal{T}}(\lambda).$$

Where does it come from: Semidefinite Relaxation

$$\psi = \max_{x \in \mathcal{X}} x^T W x, \quad \mathcal{X} = \{x : \exists(t \in \mathcal{T}, y) : x = Fy, y^T S_k y \leq t_k, k \leq K\} \quad (*)$$

- Let us denote $X = xx^T$, then $(*)$ can be rewritten equivalently:

$$\max \left\{ \text{Tr}(WX) : \begin{array}{l} \exists(t \in \mathcal{T}, Q \in \mathbf{S}_n) : X = FQF^T, \text{Tr}(S_k Q) \leq t_k, k \leq K, \\ Q \succeq 0, \text{Rank}(Q) = 1 \end{array} \right\}$$

Then $\psi \leq \bar{\psi}$ where

$$\bar{\psi} = \max_{Q \in \mathcal{Q}} \text{Tr}(\underbrace{F^T W F}_{=: V} Q), \quad \mathcal{Q} = \{Q \succeq 0 : \exists t \in \mathcal{T} : \text{Tr}(S_k Q) \leq t_k, k \leq K\} \quad (*_1)$$

Note that $(*_1)$ is a *convex program*, it is instructive to write it as *Conic Programming*.

- Conic representation of constraint $t \in \mathcal{T}$.

- Let

$$\mathcal{T}^+ = \{[t; 1] \in \mathbb{R}^{K+1} : t \in \mathcal{T}\}, \quad \mathbf{T} = \{\lambda s \in \mathbb{R}^{K+1}, \lambda \geq 0, s \in \mathcal{T}^+\}.$$

- $\mathbf{T}$ is a regular cone (since $\mathcal{T}$ is a convex compact set with a nonempty interior)
- $\mathcal{T} = \{t : [t; 1] \in \mathbf{T}\}$

The dual cone to $\mathbf{T} = \mathbb{R}_+ \times \{[t; 1] : t \in \mathcal{T}\}$ is

$$\left\{ [y; s] \in \mathbb{R}^{k+1} : s \geq \phi_{\mathcal{T}}(-y) \right\}$$

Indeed, if $\mathbf{T}_*$ is the cone dual to $\mathbf{T}$,

$$\begin{aligned} \{[y; s] \in \mathbf{T}_*\} &\Leftrightarrow \{y^T t + s\tau \geq 0 \ \forall [t; \tau] \in \mathbf{T}\} \\ &\Leftrightarrow \{y^T t + s \geq 0 \ \forall t : [t; 1] \in \mathbf{T}\} \Leftrightarrow s \geq -y^T t \ \forall t \in \mathcal{T} \Leftrightarrow s \geq \max_{t \in \mathcal{T}} [-y]^T t \\ &\Leftrightarrow \{s \geq \phi_{\mathcal{T}}(-y)\} \end{aligned}$$

Bottom line: We have

$$\overline{\Psi} = \max_{Q, t} \{ \text{Tr}(VQ) : [t; 1] \in \mathbf{T}, Q \succeq 0, \text{Tr}(S_k Q) \leq t_k \ \forall k \} \quad (*_1)$$

Question: What is the conic dual to $(*_1)$?

- To pass from $(*_1)$ to the dual problem we associate Lagrange multipliers to the constraints:

$$\underbrace{[t; 1] \in \mathbf{T}}_{[g; s] \in \mathbf{T}_*}, \quad \underbrace{Q \succeq 0}_{P \succeq 0}, \quad \underbrace{t_k - \text{Tr}(QS_k) \geq 0}_{\lambda_k \geq 0} \quad \forall k \leq K$$

- Taking inner products of the constraints with Lagrange multipliers and summing up the results, we get the consequence of the constraints of $(*_1)$:

$$g^T t + s \cdot 1 + \text{Tr}(PQ) + \sum_k \lambda_k [t_k - \text{Tr}(QS_k)] \geq 0$$

$\downarrow \uparrow$

$$\text{Tr}([\sum_k \lambda_k S_k - P] Q) - [g + \lambda]^T t \leq s, \quad [s \geq \phi_{\mathcal{T}}(-g), P \succeq 0, \lambda \geq 0]$$

- To get the dual problem, we impose on Lagrange multiplies the requirement that the body of the aggregated constraint, as a function of Q, t , is identically equal to the objective $\text{Tr}(VQ)$ of $(*_1)$:

$$\overline{\Psi} = \min_{g,s,\lambda,P} \left\{ s : \begin{array}{l} g + \lambda = 0, \quad \sum_k \lambda_k S_k - P = V \\ s \geq \phi_{\mathcal{T}}(-g), \quad P \succeq 0, \quad \lambda \geq 0 \end{array} \right\}$$

$$\Updownarrow$$

$$\overline{\Psi} = \min_{g,s,\lambda,P} \left\{ s : \begin{array}{l} g = -\lambda, \quad \sum_k \lambda_k S_k - P = V \\ s \geq \phi_{\mathcal{T}}(-g), \quad P \succeq 0, \quad \lambda \geq 0 \end{array} \right\}$$

$$= \min_{s,\lambda} \left\{ s : \begin{array}{l} \sum_k \lambda_k S_k \succeq V \\ s \geq \phi_{\mathcal{T}}(\lambda), \quad \lambda \geq 0 \end{array} \right\}$$

$$= \min_{\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) : \underbrace{\sum_k \lambda_k S_k}_{\sum_k \lambda_k S_k \succeq F^T W F} \succeq V, \quad \lambda \geq 0 \right\}$$

$$\mathcal{X} = \{x : \exists(t \in \mathcal{T}, y) : x = Fy, y^T S_k y \leq t_k, k \leq K\} \quad (*)$$

Corollary *Let $\mathcal{X}$ be the ellitope $(*)$. Then the function*

$$\Psi(H) = \max_{x \in \mathcal{X}} x^T [(B^T - A^T H)(B - H^T A)] x$$

can be upper-bounded: $\Psi(H) \leq \overline{\Psi}(H)$ where

$$\begin{aligned} \overline{\Psi}(H) &:= \min_{\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) : \lambda \geq 0, F^T [B^T - A^T H] [B - H^T A] F \preceq \sum_k \lambda_k S_k \right\} \\ &\quad \text{[Schur Complement Lemma]} \\ &= \min_{\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) : \lambda \geq 0, \left[\begin{array}{c|c} \sum_k \lambda_k S_k & F^T [B^T - A^T H] \\ \hline [B - H^T A] F & I_{\nu} \end{array} \right] \succeq 0 \right\} \end{aligned}$$

The function $\overline{\Psi}(H)$ is real-valued and convex, and is efficiently computable whenever $\phi_{\mathcal{T}}$ is so, that is, whenever $\mathcal{T}$ is computationally tractable.

Bottom line: Given matrices $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{\nu \times n}$ and an ellitope

$$\mathcal{X} = \{x : \exists(t \in \mathcal{T}, y) : x = Fy, y^T S_k y \leq t_k, k \leq K\} \quad (*)$$

contained in $\mathbb{R}^n$, consider the convex optimization problem

$$\text{Opt} = \min_{H, \lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(H^T H) : \lambda \geq 0, \left[\frac{\sum_k \lambda_k S_k}{[B - H^T A]F} \middle| \frac{F^T [B^T - A^T H]}{I_\nu} \right] \succeq 0 \right\}$$

$$\left[\phi_{\mathcal{T}}(\lambda) = \max_{t \in \mathcal{T}} t^T \lambda : \text{support function of } \mathcal{T} \right]$$

Assuming the noise ξ in observation $\omega = Ax + \sigma\xi$ zero mean with unit covariance matrix, the risk of the linear estimate $\hat{x}_{H_*}(\cdot)$ induced by the optimal solution H_* to the problem (this solution clearly exists provided $\sigma > 0$) satisfies the risk bound

$$\text{Risk}[\hat{x}_{H_*} | \mathcal{X}] \leq \sqrt{\text{Opt}}.$$

What is ahead: We are about to prove that *in the case of* $\xi \sim \mathcal{N}(0, I_m)$, $\sqrt{\text{Opt}}$ is “nearly” the same as the ideal minimax risk

$$\text{RiskOpt} = \inf_{\hat{x}(\cdot)} \text{Risk}[\hat{x} | \mathcal{X}],$$

where $\inf$ is taken w.r.t. all (not necessarily linear) estimates $\hat{x}(\cdot)$.

How it works: Inverse Heat Equation

Square plate is heated at time 0 and is left to cool; the temperature at the plate's boundary is kept 0 all the time.

Given given noisy measurements, taken along m points, of plate's temperature at time t_1 , we want to recover distribution of temperature at time a given time t_0 , $0 < t_0 < t_1$.

The model: The temperature field $u(t; p, q)$ evolves according to *Heat Equation*

$$\frac{\partial}{\partial t} u(t; p, q) = \left[\frac{\partial^2}{\partial p^2} + \frac{\partial^2}{\partial q^2} \right] u(t; p, q), \quad t \geq 0, (p, q) \in S$$

• t : time • $S = \{(p, q), -1 \leq p, q \leq 1\}$: the plate

with boundary conditions $u(t; p, q)|_{(p,q) \in \partial S} \equiv 0$.

- It is convenient to represent $u(t; p, q)$ by its expansion

$$u(t; p, q) = \sum_{k,\ell} x_{k\ell}(t) \phi_k(p) \phi_\ell(q), \quad (*)$$
$$\phi_k(s) = \begin{cases} \cos(\omega_{2i-1}s), & \omega_{2i-1} = (i - 1/2)\pi, \quad k = 2i - 1 \\ \sin(\omega_{2i}s), & \omega_{2i} = i\pi, \quad k = 2i \end{cases}$$

Note: $\phi_k(s)$ are harmonic oscillations vanishing at $s = \pm 1$.

$$\begin{aligned}
u(t; p, q) &= \sum_{k, \ell} x_{k\ell}(t) \phi_k(p) \phi_\ell(q), \\
\phi_k(s) &= \begin{cases} \cos(\omega_{2i-1}s), & \omega_{2i-1} = (i - 1/2)\pi, & k = 2i - 1 \\ \sin(\omega_{2i}s), & \omega_{2i} = i\pi, & k = 2i \end{cases}
\end{aligned} \tag{*}$$

- $\{\phi_{k\ell}(p, q) = \phi_k(p)\phi_\ell(q)\}_{k, \ell}$ form an orthonormal system in $L_2(S)$
- $\phi_{k\ell}(\cdot)$ meet the boundary conditions

$$\phi_{k\ell}(p, q) \Big|_{(p, q) \in \partial S} = 0$$

- in terms of the coefficients $x_{k\ell}(t)$, the Heat Equation becomes

$$\frac{d}{dt}x_{k\ell}(t) = -[\omega_k^2 + \omega_\ell^2]x_{k\ell}(t) \Rightarrow x_{k\ell}(t) = e^{-[\omega_k^2 + \omega_\ell^2]t}x_{k\ell}(0).$$

We select integer discretization parameter N and

- restrict (*) to terms with $1 \leq k, \ell \leq 2N - 1$
- discretize the spatial variable (p, q) to reside in the grid

$$G_N = \left\{ P_{ij} = (p_i, p_j) = \left(\frac{i}{N} - 1, \frac{j}{N} - 1 \right), 1 \leq i, j \leq 2N - 1 \right\}$$

Note: Restricting functions $\phi_{k\ell}(\cdot, \cdot)$, $1 \leq k, \ell \leq 2N - 1$ on grid G_N , we get orthogonal basis in $\mathbb{R}^{(2N-1) \times (2N-1)}$.

We arrive at the model as follows:

- The signal x underlying observation is

$$x = \{x_{k\ell} := x_{k\ell}(t_0), 1 \leq k, \ell \leq 2N - 1\} \in \mathbb{R}^{(2N-1) \times (2N-1)}$$

- The observation is

$$\omega = Ax + \sigma\xi \in \mathbb{R}^m, \quad \xi \sim \mathcal{N}(0, I_m)$$

$$[Ax]_\nu = \sum_{k,\ell=1}^{2N-1} x_{k\ell} e^{-[\omega_k^2 + \omega_\ell^2][t_1 - t_0]} \phi_k(p_{i(\nu)}) \phi_\ell(p_{j(\nu)})$$

- $(p_{i(\nu)}, p_{j(\nu)}) \in S, \quad 1 \leq \nu \leq m$: measurement points

- We want to recover the restriction Bx of $u(t_0; p, q)$ to some grid, say, square grid

$$G_K = \left\{ \left(r_i = \frac{i}{K} - 1, r_j = \frac{j}{K} - 1 \right), \quad 1 \leq i, j \leq 2K - 1 \right\} \subset S,$$

resulting in

$$[Bx]_{ij} = \sum_{k,\ell=1}^{2N-1} \phi_k(r_i) \phi_\ell(r_j) x_{k\ell}$$

- We assume that the initial distribution of temperatures $[u(0; p_i, p_j)]_{i,j=1}^{2N-1}$ satisfies $\|u\|_2 \leq R$, for some given R , implying that x resides in the ellitope, namely, the ellipsoid

$$\mathcal{X} = \left\{ \{x_{k\ell}\} \in \mathbb{R}^{(2N-1) \times (2N-1)} : \sum_{k,\ell} \left[e^{[\omega_k^2 + \omega_\ell^2]t_0} x_{k\ell} \right]^2 \leq R^2 \right\}$$

$$u(t; p_i, p_j) = \sum_{k,\ell} e^{-[\omega_k^2 + \omega_\ell^2][t-t_0]} \phi_k(p_i) \phi_\ell(p_j) x_{k\ell}$$

$$[A(x)]_\nu = \sum_{k,\ell=1}^{2N-1} x_{k\ell} e^{-[\omega_k^2 + \omega_\ell^2][t_1-t_0]} \phi_k(p_{i(\nu)}) \phi_\ell(p_{j(\nu)}) x_{k\ell}$$

Plots of $e^{-(\omega_k^2 + \omega_\ell^2)t}$ for $t = 0.02$ (red) and $t = 0.01$ (blue)

Bad news: contributions of high frequency (with large $\omega_k^2 + \omega_\ell^2$) components $x_{k\ell}(0)$ to $A(x)$ decrease exponentially fast with high decay rate as t_1 grows
 $\Rightarrow$ High frequency components $x_{k\ell}(0)$ are impossible to recover from observations at time t_1 , unless t_1 is very small (red curve, $t = 0.02$)

Good news: Contributions of high frequency components $x_{k\ell}(0)$ to $B(x)$ are very small, provided t_0 is not too small
 $\Rightarrow$ There is no need to recover well high frequency components, provided they are not huge (blue curve, $t = 0.01$)

Numerical results $N = 32$, $m = 125$, $K = 6$, $t_0 = 0.01$, $t_1 = 0.02$, $\sigma = 0.001$,

$$\mathcal{X} = \{x \in \mathbb{R}^{63 \times 63} : \|\{u(0; p_i, p_j)\}_{i,j=1}^{63}\|_2 \leq 15\}.$$

- Minimax risk of optimal linear estimate: 0.1707

63×63 grid G_{63} and $m = 125$ measurement points

B

$$\|B\|_2 = 2.128$$

$$\|B\|_\infty = 0.433$$

$\hat{B}$

$$\|B - \hat{B}\|_2 = 0.1535$$

$$\|B - \hat{B}\|_\infty = 0.0313$$

$\tilde{B}$

$$\|B - \tilde{B}\|_2 = 266.3$$

$$\|B - \tilde{B}\|_\infty = 63.63$$

Left: true signal, center: linear recovery, right: naive LS recovery

Near-optimality of linear estimates

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists y \in \mathbb{R}^N, t \in \mathcal{T} : x = Fy, y^T S_k y \leq t_k \ \forall k \leq K\}$$

Simple observation: When recovering Bx , $x \in \mathcal{X}$, from observation $\omega = Ax + \sigma\xi$, we lose nothing when assuming that the signal is

$$y \in \mathcal{Y} = \{y : \exists t \in \mathcal{T} : y^T S_k y \leq t_k \ \forall k \leq K\}$$

the observation is $\omega = [AF]y + \sigma\xi$, and the entity to be recovered is $[BF]y$.

With this transformation, families of all estimates, all linear estimates and their risks remain intact

$\Rightarrow$ *We lose nothing when assuming that F is the identity, and*

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists t \in \mathcal{T} : x^T S_k x \leq t_k \ \forall k \leq K\}$$

$\mathcal{X} = \{x \in \mathbb{R}^n : \exists t \in \mathcal{T} : x^T S_k x \leq t_k \ \forall k \leq K\} \Rightarrow$ *We can build linear estimate satisfying*
 $\text{Risk}^2[\hat{x}|\mathcal{X}] \leq \text{Opt} = \min_{H, \lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(H^T H) : \lambda \geq 0, \left[\frac{\sum_k \lambda_k S_k}{B - H^T A} \mid \frac{B^T - A^T H}{I_\nu} \right] \succeq 0 \right\}$

Our course of actions is as follows.

A. Suppose that we can specify Gaussian prior $\mathcal{N}(0, Q)$ for the signal x to be supported on $\mathcal{X}$. Then by Gauss-Markov Theorem the quantity

$$\psi(Q) = \text{Tr}(BQB^T - BQA^T[\sigma^2 I_m + AQA^T]^{-1}AQB^T)$$

would be a lower bound on RiskOpt^2 .

B. While $\mathcal{N}(0, Q)$ cannot be supported on the *bounded* set $\mathcal{X}$ for a nonzero Q , we can select Q to enforce $\eta \sim \mathcal{N}(0, Q)$ to “sit in $\mathcal{X}$ with overwhelming probability,” so that a “slightly reduced” $\psi(Q)$ will become a lower bound on RiskOpt^2 .

C. We have $\mathbb{E}_{\eta \sim \mathcal{N}(0, Q)} \{\eta^T S \eta\} = \text{Tr}(SQ)$

$\Rightarrow$ Selecting $Q \succeq 0$ according to

$$\exists t \in \mathcal{T} : \text{Tr}(QS_k) \leq t_k, k \leq K$$

we ensure that $\eta \sim \mathcal{N}(0, Q)$ “sits in $\mathcal{X}$ on average”

$\Rightarrow$ We may expect that imposing on $Q \succeq 0$ the restriction

$$\exists t \in \mathcal{T} : \text{Tr}(QS_k) \leq \rho t_k, k \leq K, \quad [\rho > 0]$$

we enforce $\eta \sim \mathcal{N}(0, Q)$ to take values in $\mathcal{X}$ with probability controlled by ρ and approaching 1 as $\rho \rightarrow +0$.

$\mathcal{X} = \{x \in \mathbb{R}^n : \exists t \in \mathcal{T} : x^T S_k x \leq t_k \ \forall k \leq K\} \Rightarrow$ *We can build linear estimate satisfying*
 $\text{Risk}^2[\hat{x}|\mathcal{X}] \leq \text{Opt} = \min_{H, \lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(H^T H) : \lambda \geq 0, \left[\frac{\sum_k \lambda_k S_k}{B - H^T A} \mid \frac{B^T - A^T H}{I_\nu} \right] \succeq 0 \right\}$

D. The above considerations give rise to parametric optimization problem

$$\text{Opt}_*(\rho) = \max_{Q \succeq 0} \{ \psi(Q) : \exists t \in \mathcal{T} : \text{Tr}(Q S_k) \leq \rho t_k \ \forall k \leq K \} \quad (*_{\rho})$$

We *can expect* that for small ρ “slightly corrected” $\text{Opt}_*(\rho)$ is a lower bound on RiskOpt^2 .

Fact: By Conic Duality , $\text{Opt}_*(\rho) \geq \rho \text{Opt}$

$\Rightarrow$ *Lower bounds on RiskOpt^2 can be expressed via Opt , resulting in*

$$\sqrt{\text{Opt}} \leq O \left(\sqrt{\ln(1/\text{RiskOpt})} \text{RiskOpt} \right)$$

provided RiskOpt is small.

From $(*_1)$ to its dual

$$\text{Opt}_{*}(1) = \max_{Q,t} \left\{ \text{Tr}(BQB^T - BQA^T[\sigma^2 I_m + AQA^T]^{-1}AQB^T) : \right. \\ \left. t \in \mathcal{T}, Q \succeq 0, \text{Tr}(QS_k) \leq t_k \forall k \right\}$$

Step 1: Passing to linear objective

$$\begin{aligned} & \overbrace{\max_{Q,t} \left\{ \text{Tr}(BQB^T - BQA^T[\sigma^2 I_m + AQA^T]^{-1}AQB^T) : t \in \mathcal{T}, Q \succeq 0, \right.}^{\text{Opt}_{*}(1)} \\ & \quad \left. \text{Tr}(QS_k) \leq t_k \forall k \right\}} \\ &= \max_{Q,t,G} \left\{ \text{Tr}(BQB^T) - \text{Tr}(G) : \begin{array}{l} G - BQA^T[\sigma^2 I_m + AQA^T]^{-1}AQB^T \succeq 0 \\ t \in \mathcal{T}, Q \succeq 0, \text{Tr}(QS_k) \leq t_k \forall k \end{array} \right\} \\ & \quad \text{[Schur Complement Lemma]} \\ &= \max_{Q,t,G} \left\{ \text{Tr}(BQB^T) - \text{Tr}(G) : \left[\begin{array}{c|c} G & BQA^T \\ \hline AQB^T & \sigma^2 I_m + AQA^T \end{array} \right] \succeq 0 \right. \\ & \quad \left. t \in \mathcal{T}, Q \succeq 0, \text{Tr}(QS_k) \leq t_k \forall k \right\} \end{aligned}$$

Step 2: Using the conic representation of $\mathcal{T}$ we get

$$\text{Opt}_{*}(1) = \max_{Q,t,G} \left\{ \text{Tr}(BQB^T) - \text{Tr}(G) : \left[\begin{array}{c|c} G & BQA^T \\ \hline AQB^T & \sigma^2 I_m + AQA^T \end{array} \right] \succeq 0 \right. \\ \left. [t; 1] \in \mathbf{T}, Q \succeq 0, \text{Tr}(QS_k) \leq t_k \forall k \right\} \quad (*)$$

Note: $(*)$ is clearly strictly feasible and solvable.

$$\text{Opt}_*(1) = \max_{Q,t,G} \left\{ \text{Tr}(BQB^T) - \text{Tr}(G) : \left[\begin{array}{c|c} G & BQA^T \\ \hline AQB^T & \sigma^2 I_m + AQA^T \end{array} \right] \succeq 0, \right. \\ \left. Q \succeq 0, [t; 1] \in \mathbf{T}, t_k - \text{Tr}(QS_k) \geq 0 \forall k \leq K \right\} \quad (*)$$

Step 3: passing from (*) to the dual problem. We start with associating Lagrange multipliers to the constraints (*recall what $\mathbf{T}_*$ is*):

$$\underbrace{\left[\begin{array}{c|c} G & BQA^T \\ \hline AQB^T & \sigma^2 I_m + AQA^T \end{array} \right] \succeq 0}_{\left[\begin{array}{c|c} U & V^T \\ \hline V & W \end{array} \right] \succeq 0}, \underbrace{[t; 1] \in \mathbf{T}}_{[g; s] \in \mathbf{T}_*}, \underbrace{Q \succeq 0}_{P \succeq 0}, \underbrace{t_k - \text{Tr}(QS_k) \geq 0}_{\lambda_k \geq 0} \quad \forall k \leq K$$

• Taking inner products of the constraints of (*) with Lagrange multipliers and summing up the results, we get the consequence of the constraints of (*):

$$\begin{aligned} & \text{Tr}(UG) + 2\text{Tr}(V^T AQB^T) + \text{Tr}(W[\sigma^2 I + AQA^T]) \\ & + g^T t + s + \text{Tr}(PQ) + \sum_k \lambda_k [t_k - \text{Tr}(QS_k)] \geq 0 \end{aligned}$$

$\Updownarrow$

$$-\text{Tr}(UG) + \text{Tr}([\sum_k \lambda_k S_k - P - A^T W A - B^T V^T A - A^T V B] Q) - [g + \lambda]^T t \leq s + \sigma^2 \text{Tr}(W)$$

$$\left[\left[\begin{array}{c|c} U & V^T \\ \hline V & W \end{array} \right] \succeq 0, s \geq \phi_{\mathcal{T}}(-g), P \succeq 0, \lambda \geq 0 \right]$$

$$-\text{Tr}(UG) + \text{Tr} \left(\left[\sum_k \lambda_k S_k - P - A^T W A - B^T V^T A - A^T V B \right] Q \right) - [g + \lambda]^T t \leq s + \sigma^2 \text{Tr}(W)$$

$$\left[\left[\begin{array}{c|c} U & V^T \\ \hline V & W \end{array} \right] \succeq 0, s \geq \phi_{\mathcal{T}}(-g), P \succeq 0, \lambda \geq 0 \right]$$

- To get the dual problem, we impose on Lagrange multipliers the requirement that the body of the aggregated – **the red** – constraint, as a function of G, Q, t , is identically equal to the objective

$$\text{Tr}(BQB^T) - \text{Tr}(G)$$

of (*), and minimize under this restriction (and the red constraints on the multipliers) the right hand side of the aggregated constraint.

By Conic Duality Theorem, we get

$$\text{Opt}_*(1) = \min_{U, V, W, g, s, \lambda, P} \left\{ s + \sigma^2 \text{Tr}(W) : \begin{array}{l} U = I_\nu, g = -\lambda \\ \sum_k \lambda_k S_k - P - A^T W A - B^T V^T A - A^T V B = B^T B \\ \left[\begin{array}{c|c} U & V^T \\ \hline V & W \end{array} \right] \succeq 0, s \geq \phi_{\mathcal{T}}(-g), P \succeq 0, \lambda \geq 0 \end{array} \right\}$$

$$\begin{aligned}
\text{Opt}_*(1) &= \min_{U,V,W,g,s,\lambda,P} \left\{ s + \sigma^2 \text{Tr}(W) : \begin{array}{l} U = I_\nu, g = -\lambda \\ \sum_k \lambda_k S_k - P - A^T W A - B^T V^T A - A^T V B = B^T B \\ \left[\begin{array}{c|c} U & V^T \\ \hline V & W \end{array} \right] \succeq 0, s \geq \phi_{\mathcal{T}}(-g), P \succeq 0, \lambda \geq 0 \end{array} \right\} \\
&= \min_{V,W,\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(W) : \underbrace{\left[\begin{array}{c|c} I_\nu & V^T \\ \hline V & W \end{array} \right] \succeq 0}_{\Leftrightarrow W \succeq V V^T}, \lambda \geq 0 \right\} \\
&= \min_{V,\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(V V^T) : \underbrace{B^T B + A^T V V^T A + B^T V^T A + A^T V B}_{=(B^T + A^T V)(B + V^T A)} \preceq \sum_k \lambda_k S_k, \lambda \geq 0 \right\} \\
&\quad \text{by Schur complement lemma} \\
&= \min_{V,\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(V V^T) : \left[\begin{array}{c|c} \sum_k \lambda_k S_k & B^T + A^T V \\ \hline B + V^T A & I_\nu \end{array} \right] \succeq 0, \lambda \geq 0 \right\} \\
&= \min_{H[:=-V],\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(H^T H) : \left[\begin{array}{c|c} \sum_k \lambda_k S_k & B^T - A^T H \\ \hline B - H^T A & I_\nu \end{array} \right] \succeq 0, \lambda \geq 0 \right\} \\
&=: \text{Opt}
\end{aligned}$$

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists t \in \mathcal{T} : x^T S_k x \leq t_k \ \forall k \leq K\} \Rightarrow \text{We can build linear estimate satisfying}$$

$$\text{Risk}^2[\hat{x}|\mathcal{X}] \leq \text{Opt} = \min_{H, \lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(H^T H) : \lambda \geq 0, \left[\begin{array}{c|c} \sum_k \lambda_k S_k & B^T - A^T H \\ \hline B - H^T A & I_\nu \end{array} \right] \succeq 0 \right\} = \text{Opt}_*(1)$$

$$\text{Opt}_*(\rho) = \max_{Q, t} \{ \psi(Q) : Q \succeq 0, t \in \mathcal{T}, \text{Tr}(Q S_k) \leq \rho t_k \ \forall k \} \quad (*_\rho)$$

$$\left[\sqrt{\psi(Q)} : \text{optimal Bayesian risk of recovering } Bx \text{ when } x \sim \mathcal{N}(0, Q) \right]$$

Easy fact: When $0 < \rho \leq 1$, one has

$$\text{Opt}_*(\rho) \geq \rho \text{Opt}_*(1) = \rho \text{Opt}.$$

Indeed, recall that

$$\psi(Q) = \text{Tr}(BQB^T - BQA^T[\sigma^2 I_m + AQA^T]^{-1}AQB^T).$$

It follows immediately that $\psi(\rho Q) \geq \rho \psi(Q)$, $0 \leq \rho \leq 1$.

Fact: If Q_ρ stems from optimal solution to $(*_\rho)$, then the probability for $\xi \sim \mathcal{N}(0, Q_\rho)$ to be outside of $\mathcal{X}$ approaches zero fast as ρ decreases:

$$\pi(\rho) := \text{Prob}_{\xi \sim \mathcal{N}(0, Q_\rho)} \{ \xi \notin \mathcal{X} \} \leq K \exp \left\{ -\frac{1 - \rho + \rho \ln(\rho)}{2\rho} \right\}$$

ρ	0.10	0.05	0.04	0.03	0.02	0.01
$\pi(\rho) \leq$	1.0e0	3.4e-2	3.1e-3	5.5e-5	1.6e-8	3.2e-18

$$K = 100$$

$\Rightarrow$ For reasonably small ρ , everything is *as if* $\mathcal{N}(0, Q_\rho)$ were supported on $\mathcal{X}$, so that, say, $0.99 \text{Opt}_*(\rho) \geq 0.99\rho \text{Opt}$ is a lower bound on RiskOpt^2 .

Theorem *Let us associate with ellitope*

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists t \in \mathcal{T} : x^T S_k x \leq t_k \ \forall k \leq K\}$$

the convex compact set

$$\mathcal{Q} = \{Q \in \mathbf{S}^n : Q \succeq 0, \exists t \in \mathcal{T} : \text{Tr}(QS_k) \leq t_k, k \leq K\},$$

and the quantity

$$M_* = \max_{Q \in \mathcal{Q}} \sqrt{\text{Tr}(BQB^T)}.$$

Then the linear estimate $\hat{x}_{H_}(\omega) = H_*^T \omega$ of Bx , $x \in \mathcal{X}$, via observation $\omega = Ax + \sigma\xi$, $\xi \sim \mathcal{N}(0, I_m)$, given by the optimal solution to the convex optimization problem*

$$\text{Opt} = \min_{H, \lambda} \left\{ \phi_{\mathcal{T}}(\lambda) + \sigma^2 \text{Tr}(H^T H) : \left[\frac{\sum_k \lambda_k S_k}{B - H^T A} \middle| \frac{B^T - A^T H}{I_\nu} \right] \succeq 0, \lambda \geq 0 \right\}$$

satisfies the risk bound

$$\text{Risk}[\hat{x}_{H_*} | \mathcal{X}] \leq \sqrt{\text{Opt}} \leq 4 \sqrt{\ln \left(\frac{6M_*^2 \sqrt{K}}{\text{Risk}_{\text{Opt}}^2[\mathcal{X}]} \right) \text{Risk}_{\text{Opt}}[\mathcal{X}]},$$

where

$$\text{Risk}_{\text{Opt}}[\mathcal{X}] = \inf_{\hat{x}(\cdot)} \sup_{x \in \mathcal{X}} \sqrt{\mathbf{E}_{\xi \sim \mathcal{N}(0, I_m)} \{ \|Bx - \hat{x}(Ax + \sigma\xi)\|_2^2 \}}$$

inf being taken with respect to all, linear and nonlinear alike, estimates $\hat{x}(\cdot)$, is the optimal minimax risk.

How it works: numerical illustration

In these experiments

- B is $n \times n$ identity matrix,
- $n \times n$ sensing matrix A is a randomly rotated matrix with singular values λ_j , $1 \leq j \leq n$, forming a geometric progression, with $\lambda_1 = 1$ and $\lambda_n = 0.01$.
- In the first experiment the signal set $\mathcal{X}_1$ is an ellipsoid:

$$\mathcal{X}_1 = \{x \in \mathbb{R}^n : \sum_{j=1}^n j^2 x_j^2 \leq 1\},$$

that is, $K = 1$, $S_1 = \sum_{j=1}^n j^2 e_j e_j^T$ (e_j are basic orths), and $\mathcal{T} = [0, 1]$.

Theoretical “suboptimality factor” in the interval [31.6, 73.7] in this experiment.

- In the second experiment, the signal set $\mathcal{X}$ is the box:

$$\mathcal{X} = \{x \in \mathbb{R}^n : |x_j| \leq 1, 1 \leq j \leq n\}$$
$$[K = n, S_k = k^2 e_k e_k^T, k = 1, \dots, K, \mathcal{T} = [0, 1]^K].$$

Theoretical “suboptimality factor” in the interval [73.2, 115.4].

Recovery on ellipsoids: risk bounds as functions of the noise level σ , dimension $n = 32$.

Left plot: upper and lower bounds of the risk; right plot: suboptimality ratios.

Recovery on ellipsoids: risk bounds as functions of problem dimension n , noise level $\sigma = 0.01$.

Left plot: upper and lower risk bounds; right plot: suboptimality ratios.

Recovery on a box: risk bounds as functions of the noise level σ , dimension $n = 32$.

Left plot: upper and lower bounds of the risk; right plot: suboptimality ratios.

Recovery on a box: risk bounds as functions of problem dimension n , noise level $\sigma = 0.01$.

Left plot: upper and lower risk bounds; right plot: suboptimality ratios.

Extension: Relative Risks

When “very large” signals are allowed, it might make sense to switch from the usual risk to its *relative version* – “ S -risk” defined as follows:

$$\text{Risk}_S[\hat{x}|\mathcal{X}] = \min \left\{ \sqrt{\tau} : \mathbf{E}_{\xi} \left\{ \|Bx - \hat{x}(Ax + \sigma\xi)\|_2^2 \right\} \leq \tau[1 + x^T S x] \forall x \in \mathcal{X} \right\}$$

S : fixed positive semidefinite “risk calibrating” matrix

Note: Setting $S = 0$ recovers the usual “plain” risk.

Results on design of near-optimal, in terms of plain risk, linear estimates extend to the case of S -risk.

Design of near optimal linear estimate $\hat{x}_{H_*}(\omega) = H_*^T \omega$ is given by an optimal solution (H_*, τ_*, λ_*) to the convex optimization problem

$$\text{Opt} = \min_{H, \tau, \lambda} \left\{ \tau : \begin{array}{c} \left[\frac{\sum_k \lambda_k S_k + \tau S}{B - H^T A} \mid \frac{B^T - A^T H}{I_\nu} \right] \succeq 0, \\ \sigma^2 \text{Tr}(H^T H) + \phi_{\mathcal{T}}(\lambda) \leq \tau, \lambda \geq 0 \end{array} \right\}$$

For the resulting estimate, it holds $\text{RiskS}[\hat{x}_{H_*} | \mathcal{X}] \leq \sqrt{\text{Opt}}$.

Near-optimality properties of the estimate $\hat{x}_{H_*}$ remain the same as in the case of plain risk:

When $\xi \sim \mathcal{N}(0, I_m)$, one has

$$\text{RiskS}[\hat{x}_{H_*} | \mathcal{X}] \leq 4 \sqrt{\ln \left(\frac{6M_*^2 \sqrt{K}}{\text{RiskS}_{\text{Opt}}^2[\mathcal{X}]} \right)} \text{RiskS}_{\text{Opt}}[\mathcal{X}],$$

where

$$M_* = \max_Q \left\{ \sqrt{\text{Tr}(BQB^T)} : Q \succeq 0, \exists t \in \mathcal{T} : \text{Tr}(QS_k) \leq t_k \ \forall k \leq K \right\}$$

and

$$\text{RiskS}_{\text{Opt}}[\mathcal{X}] = \inf_{\hat{x}(\cdot)} \text{RiskS}[\hat{x} | \mathcal{X}].$$

In the special case $\mathcal{X} = \mathbb{R}^n$, the best linear estimate is yielded by optimal solution to the convex problem

$$\text{Opt} = \min_{H, \tau} \left\{ \tau : \left[\begin{array}{c|c} \tau S & B^T - A^T H \\ \hline B - H^T A & I_\nu \end{array} \right] \succeq 0, \sigma^2 \text{Tr}(H^T H) \leq \tau \right\} \quad (*)$$

A feasible solution τ, H to $(*)$ gives rise to linear estimate $\hat{x}_H(\omega) = H^T \omega$ such that

$$\text{RiskS}[\hat{x}_H | \mathbb{R}^n] \leq \sqrt{\tau},$$

provided ξ is zero mean and with unit covariance matrix.

Proposition *Assume that $B \neq 0$ and $(*)$ is feasible. Then the problem is solvable, and its optimal solution Opt, H_* gives rise to linear estimate $\hat{x}_{H_*}(\omega) = H_*^T \omega$ with S -risk $\sqrt{\text{Opt}}$. When $\xi \sim \mathcal{N}(0, I_m)$, this estimate is minimax optimal:*

$$\text{RiskS}[\hat{x}_{H_*} | \mathbb{R}^n] = \sqrt{\text{Opt}} = \text{RiskS}_{\text{Opt}}[\mathbb{R}^n].$$

A Byproduct: accuracy of Semidefinite Relaxation

Theorem Let C be a symmetric $n \times n$ matrix and $\mathcal{X}$ be an ellitope:

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists(t \in \mathcal{T}, y) : x = Fy, y^T S_k y \leq t_k \ \forall k \leq K\}.$$

Then the efficiently computable quantity

$$\text{Opt} = \min_{\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) : \lambda \geq 0, F^T C F \preceq \sum_k \lambda_k S_k \right\} \\ \left[\phi_{\mathcal{T}}(\lambda) = \max_{t \in \mathcal{T}} \lambda^T t \right]$$

is an upper bound on

$$\text{Opt}_* = \max_{x \in \mathcal{X}} x^T C x.$$

The above bound is tight, namely,

$$\text{Opt}_* \leq \text{Opt} \leq 4 \ln(5K) \text{Opt}_*.$$

Note: Opt_* is difficult to compute within 4% accuracy when $\mathcal{X}$ is as simple as the unit box in $\mathbb{R}^n$.

Let $\mathcal{X}$ be given by quadratic inequalities:

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists t \in \mathcal{T} : x^T S_k x \leq t_k, k \leq K\} \neq \emptyset$$

[$\mathcal{T}$: nonempty convex compact set]

We have

$$\text{Opt}_* := \max_{x \in \mathcal{X}} x^T C x \leq \underbrace{\min_{\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) : \lambda \geq 0, C \preceq \sum_k \lambda_k S_k \right\}}_{=: \text{Opt}} \leq \Theta \text{Opt}_*$$

Question: What can be said about the tightness factor Θ ?

A. When $K = 1$, assuming Slater condition: $\bar{x}^T S_1 \bar{x} < t$ for some $\bar{x}$ and some $t \in \mathcal{T}$, one can set $\Theta = 1$. [the celebrated **S-Lemma**]

B. Assuming that $x^T S_k x = x_k^2$, $k \leq K = \dim x$, $\mathcal{T} = [0; 1]^K$, and C is Laplacian of a graph (i.e., off-diagonal entries in C are nonpositive and all row sums are zero), one can set $\Theta = 1.1382...$ [MAXCUT Theorem of Goemans and Williamson, 1995]

Note: Laplacian of a graph always is $\succeq 0$

C. When $C \succeq 0$ and all matrices S_k are diagonal, one can set

$$\Theta = \frac{\pi}{2} = 1.5708\dots$$

$[\frac{\pi}{2}]$ Theorem, Nesterov, 1998]

D. Assuming $\mathcal{X}$ is an ellitope (i.e., $S_k \succeq 0$, $\sum_k S_k \succ 0$ and $\mathcal{T}$ contains a positive vector), one can set $\Theta = 4 \ln(5K)$

Note: In the case of **D**, Θ indeed can be as large as $O(\ln(K))$

A byproduct of the theorem is the following useful fact:

Theorem [upper-bounding of operator norms] *Let $\|\cdot\|_x$ be a norm on $\mathbb{R}^N$ such that the unit ball $\mathcal{X}$ of the norm is an ellitope:*

$$\mathcal{X} := \{x : \|x\|_x \leq 1\} = \{x : \exists(t \in \mathcal{T}, y) : x = Py, y^T S_k y \leq t_k, k \leq K\}$$

Let, further, $\|\cdot\|$ be a norm on $\mathbb{R}^M$ such that the unit ball $\mathcal{B}^$ of the norm $\|\cdot\|_*$ conjugate to $\|\cdot\|$ is an ellitope:*

$$\begin{aligned} \mathcal{B}_* &:= \{u \in \mathbb{R}^m : u^T v \leq 1 \forall (v, \|v\| \leq 1)\} \\ &= \{u : \exists(r \in \mathcal{R}, z) : u = Qz, z^T R_\ell z \leq r_\ell, \ell \leq L\} \end{aligned}$$

Then the efficiently computable quantity

$$\text{Opt}(C) = \min_{\lambda, \mu} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\mu) : \lambda \geq 0, \mu \geq 0, \left[\frac{\sum_{\ell} \mu_{\ell} R_{\ell}}{\frac{1}{2} P^T C^T Q} \mid \frac{\frac{1}{2} Q^T C P}{\sum_k \lambda_k S_k} \right] \succeq 0 \right\}$$

$$\left[C \in \mathbb{R}^{M \times N} \right]$$

is a convex in C upper bound on the operator norm

$$\|C\|_{\|\cdot\|_x \rightarrow \|\cdot\|} = \max_x \{\|Cx\| : \|x\|_x \leq 1\}$$

of the mapping $x \mapsto Cx$, and this bound is reasonably tight:

$$\|C\|_{\|\cdot\|_x \rightarrow \|\cdot\|} \leq \text{Opt}(C) \leq 4 \ln(5(K + L)) \|C\|_{\|\cdot\|_x \rightarrow \|\cdot\|}.$$

Indeed, the operator norm in question is the maximum of a quadratic form over an ellitope:

$$\begin{aligned}\|C\|_{\|\cdot\|_x \rightarrow \|\cdot\|} &= \max_{x \in \mathcal{X}, u \in \mathcal{B}_*} [x; u]^T \left[\frac{1}{2} C^T \mid \frac{1}{2} C \right] [x; u] \\ &= \max_{[y; z] \in \mathcal{W}} [y; z]^T \left[\frac{1}{2} P^T C^T Q \mid \frac{1}{2} Q^T C P \right] [y; z]\end{aligned}$$

where $\mathcal{W}$ is the basic ellitope given by

$$\mathcal{W} = \left\{ [y; z] : \exists [t; r] \in \mathcal{T} \times \mathcal{R} : y^T S_k y \leq t_k, k \leq K, z^T R_\ell z \leq r_\ell, \ell \leq L \right\}.$$

Inside Semidefinite Relaxation

$$\text{Opt}_* = \max_{x,t} \left\{ x^T C x : x^T S_k x \leq t_k, k \leq K, t \in \mathcal{T} \right\} \quad (*)$$

Fact: Semidefinite relaxation admits an alternative description as follows:

Let us associate with () another optimization problem where instead of deterministic candidate solutions (x, t) we are looking for **random solutions** (ξ, τ) satisfying the constraints on average:*

$$\text{Opt}^+ = \max_{\xi, \tau} \left\{ \mathbf{E}\{\xi^T C \xi\} : \begin{array}{l} \mathbf{E}\{\xi^T S_k \xi\} \leq \mathbf{E}\{\tau_k\} \\ \mathbf{E}\{\tau\} \in \mathcal{T} \end{array} \right\} \quad (!)$$

Immediate observation: *Property of a random solution (ξ, τ) to be feasible for (!) depends solely on the matrix $Q = \mathbf{E}\{\xi \xi^T\}$ and the vector $t = \mathbf{E}\{\tau\}$, so that*

$$\text{Opt}^+ = \max_{Q,t} \left\{ \text{Tr}(CQ) : \begin{array}{l} \text{Tr}(S_k Q) \leq t_k \\ Q \succeq 0, t \in \mathcal{T} \end{array} \right\} = \max_{Q,t} \left\{ \text{Tr}(CQ) : \begin{array}{l} \text{Tr}(S_k Q) \leq t_k \\ Q \succeq 0, [t; 1] \in \mathbf{T} \end{array} \right\} \quad (\#)$$

$$\text{Opt}_* = \max_{x,t} \{x^T C x : \exists(t \in \mathcal{T}) : x^T S_k x \leq t_k\} \quad (*)$$

$$\text{Opt}^+ = \max_{Q,t} \{\text{Tr}(CQ) : \text{Tr}(S_k Q) \leq t_k, Q \succeq 0, [t; 1] \in \mathbf{T}\} \quad (\#)$$

Note: $(\#)$ is strictly feasible and bounded, and *the problem*

$$\text{Opt} = \min_{\lambda} \left\{ \phi_{\mathcal{T}}(\lambda) : \lambda \geq 0, C \preceq \sum_k \lambda_k S_k \right\}$$

specifying Semidefinite Relaxation upper bound on Opt is nothing but the conic dual to $(\#)$

$$\Rightarrow \text{Opt}^+ = \text{Opt}.$$

(#) suggests the following recipe for quantifying the conservatism of the upper bound Opt on Opt_* :

- Find an optimal solution Q_*, t_* to (#) and treat $Q_* \succeq 0$ as the matrix of second moments of random vector ξ (many options!)
- Random solutions (ξ, t_*) satisfy (*) “at average.” Try to “correct” them to get feasible solutions to (*) and look how “costly” this correction is in terms of the objective.

For example, in Goemans-Williamson MAXCUT and in Nesterov’s $\pi/2$ theorems, where $x^T C x$ is maximized over the unit box $\{\|x\|_\infty \leq 1\}$, one selects $\xi \sim \mathcal{N}(0, Q_*)$ and “corrects” ξ according to $\xi \mapsto \text{sign}[\xi]$.

$$\text{Opt}_* = \max_{x,t} \{x^T C x : \exists(t \in \mathcal{T}) : x^T S_k x \leq t_k\} \quad (*)$$

$$\text{Opt}^+ = \max_{Q,t} \{\text{Tr}(CQ) : \text{Tr}(S_k Q) \leq t_k, Q \succeq 0, [t; 1] \in \mathbf{T}\} \quad (\#)$$

This is how the above recipe works in the general ellitopic case:

A. Let (Q_*, t^*) be an optimal solution to $(\#)$. Set

$$\bar{C} := Q_*^{1/2} C Q_*^{1/2} = U D U^T$$

(U is orthogonal, D is diagonal).

B. Let $\xi = Q_*^{1/2} U \zeta$ with Rademacher r.v. ζ (ζ_i take values ± 1 with probability $1/2$ and are independent across i), so that $\mathbf{E}\{\xi \xi^T\} = Q_*$. Note that

$$\begin{aligned} \xi^T C \xi &= \zeta^T U^T [Q_*^{1/2} C Q_*^{1/2}] U \zeta = \zeta^T D \zeta \equiv \text{Tr}(D) \\ &= \text{Tr}(Q_*^{1/2} C Q_*^{1/2}) \equiv \text{Tr}(C Q_*) = \text{Opt}, \\ \mathbf{E}\{\xi^T S_k \xi\} &= \mathbf{E}\{\zeta^T U^T Q_*^{1/2} S_k Q_*^{1/2} U \zeta\} = \text{Tr}(U^T Q_*^{1/2} S_k Q_*^{1/2} U) \\ &= \text{Tr}(Q_*^{1/2} S_k Q_*^{1/2}) = \text{Tr}(S_k Q_*) \leq t_k^*, \quad k \leq K \end{aligned}$$

$$\xi^T C \xi \equiv \text{Opt} \quad (a) \qquad \mathbf{E}\{\xi^T S_k \xi\} \leq t_k^*, k \leq K \quad (b)$$

C. Since $S_k \succeq 0$ and ξ is “light-tail” (it comes from Rademacher random vector), simple bounds on probabilities of large deviations combine with (b) to imply that

$$\forall(\gamma \geq 0, k \leq K) : \text{Prob}\{\xi : \xi^T S_k \xi > \gamma t_k^*\} \leq O(1) \exp\{-O(1)\gamma\}$$

$\Rightarrow$ with $\gamma_* = O(1) \ln(K+1)$, there exists a realization $\hat{\xi}$ of ξ such that

$$\hat{\xi}^T S_k \hat{\xi} \leq \gamma_* t_k^*, \quad k \leq K$$

$\Rightarrow (\xi^* = \frac{\hat{\xi}}{\sqrt{\gamma_*}}, t^*)$ is feasible for

$$\text{Opt}_* = \max_{x,t} \left\{ x^T C x : \exists(t \in \mathcal{T}) : x^T S_k x \leq t_k \right\} \quad (*)$$

$\Rightarrow \text{Opt}_* \geq \gamma_*^{-1} \hat{\xi}^T C \hat{\xi} = \gamma_*^{-1} \text{Opt}$ (look at (a)!)

A simple bounds on deviation probabilities stems from the following

Mini-Lemma *Let P be positive semidefinite $N \times N$ matrix with trace ≤ 1 and ζ be N -dimensional Rademacher random vector. Then*

$$\mathbb{E} \left\{ \exp \left\{ \zeta^T P \zeta / 3 \right\} \right\} \leq \sqrt{3} \Rightarrow \text{Prob} \{ \zeta^T P \zeta > \gamma \} \leq \sqrt{3} e^{-\gamma/3}$$

Mini-Lemma $\Rightarrow$ bounds: We have

$$\xi^T S_k \xi = \zeta^T \underbrace{U^T Q_*^{1/2} S_k Q_*^{1/2} U}_{=: t_k^* P_k} \zeta$$

with $\text{Tr}(P_k) = \text{Tr}(Q_*^{1/2} S_k Q_*^{1/2}) / t_k^* = \text{Tr}(S_k Q_*) / t_k^* \leq 1$

$\Rightarrow$ [Mini-Lemma] $\text{Prob} \{ \zeta^T P_k \zeta > 3\rho \} \leq \sqrt{3} e^{-\rho}$

$\Rightarrow \text{Prob} \{ \xi^T S_k \xi > \gamma t_k^* \} = \text{Prob} \{ \zeta^T P_k \zeta > \gamma \} \leq \sqrt{3} e^{-\gamma/3}.$

Proof of the lemma: Let $P = \sum_i \sigma_i f_i f_i^T$ be the eigenvalue decomposition of P , so that $f_i^T f_i = 1$, $\sigma_i \geq 0$, and $\sum_i \sigma_i \leq 1$. The function

$$F(\sigma_1, \dots, \sigma_N) = \mathbf{E} \left\{ e^{\frac{1}{3} \sum_i \sigma_i \zeta^T f_i f_i^T \zeta} \right\}$$

is convex on the simplex $\{\sigma \geq 0, \sum_i \sigma_i \leq 1\}$ and thus attains its maximum over the simplex at a vertex, implying that for some $h = f_i$, $h^T h = 1$, it holds

$$\mathbf{E} \left\{ e^{\frac{1}{3} \zeta^T P \zeta} \right\} \leq \mathbf{E} \left\{ e^{\frac{1}{3} (h^T \zeta)^2} \right\}.$$

Let $\xi \sim \mathcal{N}(0, 1)$ be independent of ζ . We have

$$\begin{aligned} \mathbf{E}_\zeta \left\{ e^{\frac{1}{3} (h^T \zeta)^2} \right\} &= \mathbf{E}_\zeta \left\{ \mathbf{E}_\xi \left\{ \exp\{[\sqrt{2/3} h^T \zeta] \xi\} \right\} \right\} = \mathbf{E}_\xi \left\{ \mathbf{E}_\zeta \left\{ \exp\{[\sqrt{2/3} h^T \zeta] \xi\} \right\} \right\} \\ &= \mathbf{E}_\xi \left\{ \prod_{s=1}^N \mathbf{E}_\zeta \left\{ \exp\{\sqrt{2/3} \xi h_s \zeta_s\} \right\} \right\} \\ &= \mathbf{E}_\xi \left\{ \prod_{s=1}^N \cosh(\sqrt{2/3} \xi h_s) \right\} \leq \mathbf{E}_\xi \left\{ \prod_{s=1}^N \exp\{\xi^2 h_s^2 / 3\} \right\} \\ &= \mathbf{E}_\xi \left\{ \exp\{\xi^2 / 3\} \right\} = \sqrt{3} \quad \square \end{aligned}$$

Extensions

So far, we have considered a problem of recovering Bx from observation

$$\omega = Ax + \eta$$

where

- x is unknown signal known to belong to a given basic ellitope

$$\mathcal{X} = \{x \in \mathbb{R}^n : \exists t \in \mathcal{T} : x^T S_k x \leq t_k, k \leq K\}$$

- $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{\nu \times n}$ are given matrices
- $\eta \sim \mathcal{N}(0, \sigma^2 I_m)$ is observation noise

item (squared) risk of a candidate estimate is the worst-case, over $x \in \mathcal{X}$, expected squared $\|\cdot\|_2$ -norm of recovery error:

$$\text{Risk}^2[\hat{x}|\mathcal{X}] = \sup_{x \in \mathcal{X}} \mathbf{E} \left\{ \|Bx - \hat{x}(Ax + \eta)\|_2^2 \right\}.$$

We are about to extend our results to the situation where

- Noise η not necessary is zero mean Gaussian; we allow the distribution P of noise to be unknown in advance and to depend on signal x .

Assumption: *The matrix of second moments of P admits a known upper bound:*

$$\mathbb{V}[P] := \mathbf{E}_{\eta \sim P} \{\eta \eta^T\} \preceq Q$$

for a given $Q \succ 0$

- We measure recovering error in a given norm $\|\cdot\|$, not necessarily the Euclidean one, and define risk of a candidate estimate $\hat{x}(\cdot)$ as

$$\text{Risk}_{Q, \|\cdot\|}[\hat{x}|\mathcal{X}] = \sup_{x \in \mathcal{X}} \sup_{P: \mathbb{V}[P] \preceq Q} \mathbf{E}_{\eta \sim P} \{\|Bx - \hat{x}(Ax + \eta)\|\}$$

Assumption: *The unit ball $\mathcal{B}_*$ of the norm conjugate to $\|\cdot\|$ is an ellitope:*

$$\|u\| = \max_{h \in \mathcal{B}^*} h^T u,$$

$$\mathcal{B}_* = \{h : \exists(y \in \mathbb{R}^M, r \in \mathcal{R}) : h = Fy, y^T R_\ell y \preceq r_\ell \ \forall \ell \leq L\}$$

Building linear estimate

We have

$$\begin{aligned}\text{Risk}_{Q, \|\cdot\|}[\hat{x}_H | \mathcal{X}] &= \sup_{x \in \mathcal{X}} \sup_{P: \text{Var}[P] \preceq Q} \mathbf{E}_{\eta \sim Q} \left\{ \|Bx - H^T[Ax + \eta]\| \right\} \\ &\leq \sup_{x \in \mathcal{X}} \sup_{P: \text{Var}[P] \preceq Q} \mathbf{E}_{\eta \sim Q} \left\{ \|[B - H^T A]x\| + \|H^T \eta\| \right\} \\ &\leq \Phi(H) + \Psi_Q[H], \\ \left[\Phi(H) = \max_{x \in \mathcal{X}} \|[B - H^T A]x\|, \quad \Psi_Q(H) = \sup_{P: \text{Var}[P] \preceq Q} \mathbf{E}_{\eta \sim P} \left\{ \|H^T \eta\| \right\} \right]\end{aligned}$$

Next,

$$\mathcal{B}_* = \{u = My : y \in \mathcal{Y}\}, \mathcal{Y} = \{y : \exists r \in \mathcal{R} : y^T R_\ell y \leq r_\ell \ \forall \ell \leq L\}$$

whence

$$\begin{aligned} \Phi(H) &:= \max_{x \in \mathcal{X}} \| [B - H^T A] x \| \\ &= \max_{[u; x] \in \mathcal{B}_* \times \mathcal{X}} [u; x]^T \left[\frac{\frac{1}{2}[B^T - A^T H]}{\left| \frac{1}{2}[B^T - A^T H] \right|} \middle| \frac{\frac{1}{2}[B - H^T A]}{\left| \frac{1}{2}[B - H^T A] \right|} \right] [u; x] \\ &= \max_{[y; x] \in \mathcal{Y} \times \mathcal{X}} [y; x]^T \left[\frac{\frac{1}{2}F^T[B - H^T A]}{\left| \frac{1}{2}[B^T - A^T H]Fy \right|} \middle| \frac{\frac{1}{2}F^T[B - H^T A]}{\left| \frac{1}{2}[B^T - A^T H]Fy \right|} \right] [y; x] \\ &\quad \text{[semidefinite relaxation; note that } \mathcal{Y} \times \mathcal{X} \text{ is an ellitope]} \end{aligned}$$

Thus,

$$\begin{aligned} \Phi(H) &\leq \overline{\Phi}(H) := \min_{\lambda, \mu} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\mu) : \left[\frac{\sum_{\ell} \mu_{\ell} R_{\ell}}{\left| \frac{1}{2}[A^T H - B^T]F \right|} \middle| \frac{\frac{1}{2}F^T[H^T A - B]}{\sum_k \lambda_k S_k} \right] \succeq 0 \right\} \\ &\quad \left[\phi_{\mathcal{T}}(\lambda) = \max_{t \in \mathcal{T}} \lambda^T t, \phi_{\mathcal{R}}(\mu) = \max_{r \in \mathcal{R}} \mu^T r \right] \end{aligned}$$

Lemma *One has*

$$\Psi_Q(H) \leq \bar{\Psi}_Q(H) := \min_{\Theta, \kappa} \left\{ \text{Tr}(Q\Theta) + \phi_{\mathcal{R}}(\kappa) : \kappa \geq 0, \underbrace{\left[\begin{array}{c|c} \sum_{\ell} \kappa_{\ell} R_{\ell} & \frac{1}{2} F^T H^T \\ \hline \frac{1}{2} H F & \Theta \end{array} \right]}_{(*)} \succeq 0 \right\}$$

Indeed, let (κ, Θ) be feasible for the right hand side problem, and let $\forall[P] \preceq Q$

$$\begin{aligned} \|H^T \xi\| &= \max_{u \in \mathcal{B}_*} [-u^T H^T \xi] = \max_{y \in \mathcal{Y}} [-y^T F^T H^T \xi] \\ &\leq \max_{y \in \mathcal{Y}} [y^T [\sum_{\ell} \kappa_{\ell} R_{\ell}] y + \xi^T \Theta \xi] \quad [\text{by } (*)] \\ &= \max_{r \in \mathcal{R}, y} \{ y^T [\sum_{\ell} \kappa_{\ell} R_{\ell}] y + \xi^T \Theta \xi : y^T R_{\ell} y \leq r_{\ell}, \ell \leq L \} \\ &\leq \max_{r \in \mathcal{R}} \{ \sum_{\ell} \kappa_{\ell} r_{\ell} + \xi^T \Theta \xi \} \\ &\leq \phi_{\mathcal{R}}(\kappa) + \xi^T \Theta \xi = \phi_{\mathcal{R}}(\kappa) + \text{Tr}(\Theta [\xi \xi^T]). \end{aligned}$$

Taking expectation in ξ , the conclusion of Lemma follows. □

Illustration: When $\|\cdot\| = \|\cdot\|_p$, $p \in [1, 2]$, Lemma states that whenever $\forall[P] \preceq Q$, one has

$$\mathbf{E}_{\eta \sim P} \{ \|H^T \eta\|_p \} \leq \left\| \left[\|\text{Col}_1[Q^{1/2} H]\|_2; \dots; \|\text{Col}_{\nu}[Q^{1/2} H]\|_2 \right] \right\|_p$$

Summary: Consider convex optimization problem

$$\text{Opt} = \min_{H, \lambda, \mu, \kappa, \Theta} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\mu) + \phi_{\mathcal{R}}(\kappa) + \text{Tr}(Q\Theta) : \begin{array}{l} \left[\begin{array}{c|c} \sum_{\ell} \mu_{\ell} R_{\ell} & \frac{1}{2} F^T [H^T A - B] \\ \hline \frac{1}{2} [A^T H - B^T] F & \sum_k \lambda_k S_k \end{array} \right] \succeq 0 \\ \left[\begin{array}{c|c} \sum_{\ell} \kappa_{\ell} R_{\ell} & \frac{1}{2} F^T H^T \\ \hline \frac{1}{2} H F & \Theta \end{array} \right] \succeq 0 \\ \lambda \geq 0, \mu \geq 0, \kappa \geq 0 \end{array} \right\}$$

The problem is solvable, and the H -component H_* of its optimal solution yields linear estimate

$$\hat{x}_{H_*}(\omega) = H_*^T \omega$$

such that

$$\text{Risk}_{Q, \|\cdot\|}[\hat{x}_{H_*} | \mathcal{X}] \leq \text{Opt}.$$

In the case of zero mean Gaussian observation noise, the estimate $\hat{x}_{H_*}$ is near-optimal.

Theorem *We have*

$$\text{Risk}_{Q, \|\cdot\|}[\hat{x}_{H_*} | \mathcal{X}] \leq \text{Opt} \leq C \sqrt{\ln(2L) \ln \left(\frac{2M_*^2 K}{\text{RiskOpt}_{Q, \|\cdot\|}^2} \right)} \text{RiskOpt}_{Q, \|\cdot\|}.$$

Here

- C is a positive absolute constant,

- $\text{RiskOpt}_{Q, \|\cdot\|} = \inf_{\hat{x}(\cdot)} \left[\sup_{x \in \mathcal{X}} \mathbf{E}_{\eta \sim \mathcal{N}(0, Q)} \{ \|Bx - \hat{x}(Ax + \eta)\| \} \right],$

the infimum being taken over all estimates, is the minimax optimal $\|\cdot\|$ -risk corresponding to $\mathcal{N}(0, Q)$ observation noise,

- K, L are the “sizes” of ellitopes $\mathcal{X}$ and $\mathcal{B}_* := F\mathcal{Y}$:

$$\mathcal{X} = \{x : \exists t \in \mathcal{T} : x^T S_k x \leq t_k, k \leq K\}, \quad \mathcal{Y} = \{y : \exists r \in \mathcal{R} : y^T R_\ell y \leq r_\ell, \ell \leq L\},$$

- $M_*^2 = \max_W \left\{ \mathbf{E}_{\eta \sim \mathcal{N}(0, I_n)} \|BW^{1/2}\eta\|^2 : W \in \mathcal{Q} \right\},$ *where*

$$\mathcal{Q} := \{W : W \succeq 0, \exists t \in \mathcal{T} : \text{Tr}(S_k W) \leq t_k, k \leq K\}$$

A key fact underlying near-optimality is important by its own right:

Proposition *When the unit ball $\mathcal{B}_*$ of the norm $\|\cdot\|_*$ conjugate to $\|\cdot\|$ is an ellitope:*

$$\mathcal{B}_* = \{u : \exists r \in \mathcal{R}, y : u = Fy, y^T R_\ell y \leq r_\ell, \ell \leq L\}$$

and $\eta \sim \mathcal{N}(0, Q)$, the upper bound

$$\mathbf{E}_\eta\{\|H^T \eta\|\} \leq \bar{\Psi}_Q(H) := \min_{\Theta, \varkappa} \left\{ \text{Tr}(Q\Theta) + \phi_{\mathcal{R}}(\varkappa) : \left[\begin{array}{c|c} \sum_\ell \varkappa_\ell R_\ell & \frac{1}{2} F^T H^T \\ \hline \frac{1}{2} H F & \Theta \end{array} \right] \succeq 0 \right\}$$

is tight:

$$\mathbf{E}_\eta\{\|H^T \eta\|\} \leq \bar{\Psi}_Q(H) \leq O(1) \sqrt{\ln(2L)} \mathbf{E}_\eta\{\|H^T \eta\|\}.$$

Variation: Recovery of partially random signals

So far, we have considered the problem of recovering the image Bx of unknown *deterministic* signal x known to belong to a given signal set $\mathcal{X}$ from noisy observations

$$\omega = Ax + \xi$$

of linear image of the signal.

In some applications, it makes sense to consider similar problem when the signal has a *random* component.

Example: Kalman Filter. Consider linear dynamical system

$$\begin{aligned}y_0 &= \zeta_0 \\y_{t+1} &= P_t y_t + u_t + \zeta_t, \quad t = 1, \dots, T \\ \omega_t &= C_t y_t + \xi_t\end{aligned}$$

- $y_t \in \mathbb{R}^n$: states
- $\omega_t \in \mathbb{R}^m$: observations
- ξ_t : random observation noise
- u_t : controls
- ζ_t : random “process noise”
- P_t, C_t : known matrices.

We want is to recover from observations $\omega_1, \dots, \omega_T$ linear image

$$z := R[y_0; \dots; y_{T+1}]$$

of the state trajectory, e.g., y_T (“filtering”) or y_{T+1} (“forecast”), etc

Note: In the classical Kalman Filter,

- $\zeta_0, \dots, \zeta_T$ are independent of each other zero mean Gaussian
- $\xi_1, \dots, \xi_T$ are independent of each other and of ζ_t ’s zero mean Gaussian
- $u_1, \dots, u_T$ are deterministic and known (reduces to the case when $u_t \equiv 0$)

- When modeling the situation as an estimation problem, we can use state equation to express the states y_t as known linear functions of controls u_t and process noises ζ_t , thus arriving at the model

$$\omega = A[u; \zeta] + \xi, \quad z = B[u; \zeta]$$

$$[u = [u_1; \dots; u_T], \zeta = [\zeta_0; \dots; \zeta_T], \xi = [\xi_1; \dots; \xi_T]]$$

- When quantifying the performance of a candidate estimate $\hat{x}(\omega)$, one can consider risk of the form

$$\text{Risk}[\hat{x}] = \sup_u \mathbf{E}_{\xi, \zeta} \{ \|B[u; \zeta] - \hat{x}(A[u; \zeta] + \xi)\| \}.$$

Situation: We observe noisy linear image

$$\omega = A[u; \zeta] + \xi = A_d u + A_s \zeta + \xi$$

of “signal” $x = [u; \zeta]$ with deterministic component u and stochastic component ζ . We assume that

- u is “uncertain-but-bounded” – is known to belong to a given set $\mathcal{U}$
- ζ and ξ have partially known distributions.

Specifically, for given $Q_\zeta \succ 0, Q_\xi \succ 0$ it holds

$$\mathbb{V}[\xi] = \mathbf{E}\{\xi\xi^T\} \preceq Q_\xi, \quad \mathbb{V}[\zeta] = \mathbf{E}\{\zeta\zeta^T\} \preceq Q_\zeta$$

Given matrix $B = [B_d, B_s]$ and a norm $\|\cdot\|$ on the image space of B , we want to recover $B[u; \zeta] = B_d u + B_s \zeta$, quantifying the recovery error in $\|\cdot\|$. The performance of a candidate estimate $\hat{x}(\cdot)$ is quantified by

$$\text{Risk}[\hat{x}] = \sup_{u \in \mathcal{U}} \sup_{P \in \mathcal{P}} \mathbf{E}_{[\xi; \zeta] \sim P} \{ \|B[u; \zeta] - \hat{x}(A[u; \zeta] + \xi)\| \}$$

$[\mathcal{P} : \text{probability distributions such that } \mathbf{E}_{[\xi; \zeta] \sim P} \{\xi\xi^T\} \preceq Q_\xi, \mathbf{E}_{[\xi; \zeta] \sim P} \{\zeta\zeta^T\} \preceq Q_\zeta]$

Goal: Build “presumably good” *linear* estimate $\hat{x}_H(\omega) = H^T \omega$.

Assumption: $\mathcal{U}$ is a basic ellitope, and the unit ball of the norm $\|\cdot\|_*$ dual to $\|\cdot\|$ is an ellitope:

$$\begin{aligned}\mathcal{U} &= \{u : \exists t \in \mathcal{T} : u^T S_k u \leq t_k, k \leq K\} \\ \{v : \|v\|_* \leq 1\} &= \{v : \exists r \in \mathcal{R}, w : v = Mw, w^T R_\ell w \leq r_\ell, \ell \leq L\}\end{aligned}$$

- For a candidate linear estimate $\hat{x}_H(\omega) = H^T \omega$, $u \in \mathcal{U}$, and a distribution P of $[\xi; \zeta]$ satisfying the bounds on the matrices of second moments of ξ and ζ we have

$$\begin{aligned}\mathbf{E}_{[\xi, \zeta] \sim P} \left\{ \|B_d u + B_s \zeta - H^T [A_d u + A_s \zeta + \xi]\| \right\} \\ \leq \|B_d - H^T A_d\| u + \mathbf{E}_{[\xi; \zeta] \sim P} \left\{ \|H^T \xi\| \right\} + \mathbf{E}_{[\xi; \zeta] \sim P} \left\{ \|[B_s - H^T A_s] \zeta\| \right\}\end{aligned}$$

As we know,

$$\begin{aligned}u \in \mathcal{U} \Rightarrow \| [B_d - H^T A_d] u \| &\leq \min_{\lambda \geq 0, \nu \geq 0} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\nu) : \left[\frac{\sum_{\ell} \nu_{\ell} R_{\ell}}{\frac{1}{2} [B_d^T - A_d^T H]} \middle| \frac{\frac{1}{2} M^T [B_d - H^T A_d^T]}{\sum_k \lambda_k S_k} \right] \succeq 0 \right\} \\ \forall [\xi] \preceq Q_{\xi} \Rightarrow \mathbf{E}_{\xi} \left\{ \|H^T \xi\| \right\} &\leq \min_{\mu \geq 0, G} \left\{ \text{Tr}(G Q_{\xi}) + \phi_{\mathcal{R}}(\mu) : \left[\frac{G}{\frac{1}{2} M^T H^T} \middle| \frac{\frac{1}{2} H M}{\sum_{\ell} \mu_{\ell} R_{\ell}} \right] \succeq 0 \right\} \\ \forall [\zeta] \preceq Q_{\zeta} \Rightarrow \mathbf{E}_{\zeta} \left\{ \|[B_s - H^T A_s] \zeta\| \right\} &\leq \min_{\mu \geq 0, G} \left\{ \begin{aligned} &\text{Tr}(G Q_{\zeta}) + \phi_{\mathcal{R}}(\mu) : \\ &\left[\frac{G}{\frac{1}{2} M^T [B_s - H^T A_s]} \middle| \frac{\frac{1}{2} [B_s^T - A_s^T H] M}{\sum_{\ell} \mu_{\ell} R_{\ell}} \right] \succeq 0 \end{aligned} \right\}\end{aligned}$$

Bottom line: *In the situation, consider the convex optimization problem*

$$\text{Opt} = \min_{\substack{H, \lambda, \nu, \\ \mu, \mu', G, G'}} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\nu) + \phi_{\mathcal{R}}(\mu) + \phi_{\mathcal{R}}(\mu') + \text{Tr}(Q_{\xi}G) + \text{Tr}(Q_{\zeta}G') : \right. \\ \left. \lambda \geq 0, \nu \geq 0, \mu \geq 0, \mu' \geq 0, \left[\frac{\sum_{\ell} \nu_{\ell} R_{\ell}}{\frac{1}{2}[B_d^T - A_d^T H]} \mid \frac{\frac{1}{2}M^T[B_d - H^T A_d^T]}{\sum_k \lambda_k S_k} \right] \succeq 0 \right. \\ \left. \left[\frac{G}{\frac{1}{2}M^T H^T} \mid \frac{\frac{1}{2}HM}{\sum_{\ell} \mu_{\ell} R_{\ell}} \right] \succeq 0, \left[\frac{G'}{\frac{1}{2}M^T[B_s - H^T A_s]} \mid \frac{\frac{1}{2}[B_s^T - A_s^T H]M}{\sum_{\ell} \mu'_{\ell} R_{\ell}} \right] \succeq 0 \right\}$$

The problem is efficiently solvable, and the H -component H_ of its optimal solution gives rise to linear estimate $\hat{x}_{H_*}(\Omega) = H_*^T \omega$ such that*

$$\text{Risk}[\hat{x}_{H_*}] \leq \text{Opt}.$$

How it works

System: Discretized pendulum $\ddot{x} = -\dot{x} - \kappa x$:

$$\begin{bmatrix} x_{t+1} \\ v_{t+1} \end{bmatrix} = \begin{bmatrix} 0.9990 & 0.0951 \\ -0.0190 & 0.9039 \end{bmatrix} \begin{bmatrix} x_t \\ v_t \end{bmatrix} + (u_t + \zeta_t) \begin{bmatrix} 0.0048 \\ 0.0951 \end{bmatrix}, \quad 1 \leq t \leq 128$$

$$\omega_t = x_t$$

$$\left[\begin{bmatrix} x_1 \\ v_1 \end{bmatrix} \right] \sim \mathcal{N}(0, I), \quad \zeta_t \sim \mathcal{N}(0, 0.05^2), \quad \xi \sim \mathcal{N}(0, 0.05^2), \quad |u_t| \leq 0.1$$

sample trajectory and forecasts

errors/mean errors/error bounds

Recovery under uncertain-but-bounded noise

So far, we have considered recovering Bx , $x \in \mathcal{X}$, from observation

$$\omega = Ax + \eta$$

affected by *random* noise η . Let us now assume that η is “uncertain-but-bounded:” all we know is that

$$\eta \in \mathcal{N}$$

with a given convex and compact set $\mathcal{H}$.

Risk: In the case in question, natural definition of risk of a candidate estimate $\hat{x}(\cdot)$ is

$$\text{Risk}_{\mathcal{H}, \|\cdot\|}[\hat{x}(\cdot)|\mathcal{X}] = \sup_{x \in \mathcal{X}, \eta \in \mathcal{H}} \|Bx - \hat{x}(Ax + \eta)\|.$$

Observation: *Signal recovery under uncertain-but-bounded noise reduces to the situation **where there is no observation noise at all**.*

Indeed, let us treat as the signal underlying observation the pair $z = [x; \eta] \in \mathcal{Z} := \mathcal{X} \times \mathcal{H}$ and replace A with $\bar{A} = [A, I]$ and B with $\bar{B} = [B, 0]$, so that

$$\omega = \bar{A}[x; \eta] \quad \& \quad Bx = \bar{B}[x; \eta],$$

thus reducing signal recovery to recovering $\bar{B}z$, $z \in \mathcal{Z}$, from noiseless observation $\bar{A}z$.

From now on, let us focus on the problem of recovering the image $Bx \in \mathbb{R}^\nu$ of unknown signal $x \in \mathbb{R}^n$ known to belong to signal set $\mathcal{X} \subset \mathbb{R}^n$ via observation

$$\omega = Ax \in \mathbb{R}^m.$$

Given norm $\|\cdot\|$ on $\mathbb{R}^\nu$, we quantify the performance of an estimate $\hat{x}(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^n$ by its $\|\cdot\|$ -risk

$$\text{Risk}_{\|\cdot\|}[\hat{x}|\mathcal{X}] = \sup_{x \in \mathcal{X}} \|Bx - \hat{x}(Ax)\|.$$

Observation: Assuming that $\mathcal{X}$ is computationally tractable convex compact set and $\|\cdot\|$ is computationally tractable, it is easy to build an efficiently computable *optimal within factor 2 nonlinear* estimate:

Given ω , let us solve the convex feasibility problem

$$\text{Find } y \in \mathcal{Y}[\omega] := \{y \in \mathcal{X} : Ay = \omega\}.$$

and take, as $\hat{x}(\omega)$, the vector By , where y is (any) solution to the feasibility problem.

Note: When ω stems from a signal $x \in \mathcal{X}$, the set $\mathcal{Y}[\omega]$ contains x
 $\Rightarrow \hat{x}(\cdot)$ is well defined

$$x \in \mathcal{X}, \omega = Ax \Rightarrow \hat{x}(\omega) = By \quad [y \in \mathcal{Y}[\omega] = \{y \in \mathcal{X} : Ay = \omega\}]$$

Performance analysis: Denote

$$\begin{aligned} \mathfrak{R} &= \max_{y,z} \left\{ \frac{1}{2} \|B[y - z]\| : y, z \in \mathcal{X}, A[y - z] = 0 \right\} \\ &= \frac{1}{2} \|B[y_* - z_*]\| \quad [y_*, z_* \in \mathcal{X}, A[y_* - z_*] = 0] \end{aligned}$$

Claim A: For every estimate $\tilde{x}(\cdot)$ it holds $\text{Risk}_{\|\cdot\|}[\tilde{x}|\mathcal{X}] \geq \mathfrak{R}$.

Indeed, the observation $\omega = Ay_* = Az_*$ stems from both y_* and z_* , whence the $\|\cdot\|$ -risk of every estimate is at least $\frac{1}{2} \|By_* - Bz_*\| = \mathfrak{R}$.

Claim B: One has $\text{Risk}_{\|\cdot\|}[\hat{x}|\mathcal{X}] \leq 2\mathfrak{R}$.

Indeed, let $\omega = Ax$ with $x \in \mathcal{X}$, and let $\hat{x}(\omega) = B\hat{y}$ with $\hat{y} \in \mathcal{Y}[\omega]$. Then both $x, \hat{y}$ belong to $\mathcal{Y}[\omega]$

$$\Rightarrow \frac{1}{2} \|B[x - \hat{y}]\| \leq \mathfrak{R}.$$

Question: We have built optimal, within factor 2, estimate. How to upper-bound its $\|\cdot\|$ -risk?

Let $\mathcal{X}$ and the unit ball $\mathcal{B}_*$ of the norm $\|\cdot\|_*$ *conjugate* to $\|\cdot\|$ be ellitopes:

$$\begin{aligned}\mathcal{X} &= \{x = Py : y \in \mathcal{Y} := \{y : \exists t \in \mathcal{T} : y^T S_k y \leq t_k, k \leq K\}\} \\ \mathcal{B}_* &= \{u = Qv : v \in \mathcal{V} := \{v : \exists r \in \mathcal{R} : v^T R_\ell v \leq r_\ell, \ell \leq L\}\}\end{aligned}$$

Then the $\|\cdot\|$ -risk of the optimal, within factor 2, efficiently computable nonlinear estimate $\hat{x}(\cdot)$ can be tightly lower- and upper-bounded as follows.

- Assume that $\text{Ker}A \cap \mathcal{X} \neq \{0\}$ (otherwise the risk is zero), the set $\mathcal{X}_A = \{x \in \mathcal{X} : Ax = 0\}$ is an ellitope:

$$\mathcal{X}_A = \{x = Fw, w \in \mathcal{W} := \{w : \exists t \in \mathcal{T} : w^T T_k w \leq t_k, k \leq K\}\}$$

Indeed, setting $E = \{y : APy = 0\}$, the set

$$\mathcal{Y}_A = \{y \in E : \exists t \in \mathcal{T} : y^T S_k y \leq t_k, k \leq K\}$$

is a basic ellitope in some $\mathbb{R}^{N'} \Rightarrow \mathcal{X}_A = \{Py : y \in \mathcal{Y}_A\}$ is an ellitope.

- Consequently,

$$\begin{aligned}
\mathfrak{R} &:= \max_{x', x'' \in \mathcal{X}} \left\{ \frac{1}{2} \|B[x' - x'']\| : A[x' - x''] = 0 \right\} \\
&= \max_{x \in \mathcal{X}_A} \|Bx\| = \max_{w \in \mathcal{W}} \|BFw\| \\
&= \|BF\|_{\|\cdot\|_w \rightarrow \|\cdot\|} \quad [\|\cdot\|_w: \text{norm with the unit ball } \mathcal{W}],
\end{aligned}$$

and

$$\mathfrak{R} \leq \text{Opt} \leq 4 \ln(5[K + L])\mathfrak{R}$$

where the efficiently computable quantity Opt is given by

$$\text{Opt} = \min_{\lambda, \mu} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\mu) : \begin{array}{c} \lambda \geq 0, \mu \geq 0 \\ \left[\begin{array}{c|c} \sum_{\ell} \mu_{\ell} R_{\ell} & \frac{1}{2} Q^T BF \\ \hline \frac{1}{2} F^T B^T Q & \sum_k \lambda_k T_k \end{array} \right] \succeq 0 \end{array} \right\}.$$

$\Rightarrow$ The optimal $\|\cdot\|$ -risk is $\geq \mathfrak{R} \geq \frac{\text{Opt}}{4 \ln(5[K+L])}$, and

$$\text{Risk}_{\|\cdot\|}[\hat{x}|\mathcal{X}] \leq 2\mathfrak{R} \leq 2\text{Opt}.$$

In fact, under mild assumptions a *linear* estimate is near-optimal:

Theorem Consider the problem of recovering Bx in $\|\cdot\|$, $x \in \mathcal{X}$, via observation $\omega = Ax$. Let the signal set $\mathcal{X}$ be an ellitope, and the unit ball $\mathcal{B}_*$ of the norm *conjugate* to $\|\cdot\|$ be a *basic* ellitope:

$$\mathcal{X} = \left\{ x = Py : y \in \mathcal{Y} := \{y : \exists t \in \mathcal{T} : y^T S_k y \leq t_k, k \leq K\} \right\},$$

$$\mathcal{B}_* = \{u : \exists r \in \mathcal{R} : u^T R_\ell u \leq r_\ell, \ell \leq L\}$$

Then the linear estimate $\hat{x}(\omega) = H_*^T \omega$ yielded by the H -component of optimal solution to the efficiently solvable optimization problem

$$\text{Opt} = \min_{\lambda, \mu, H} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\mu) : \left[\begin{array}{c|c} \sum_{\ell} \mu_{\ell} R_{\ell} & \frac{1}{2}[B - H^T A]P \\ \hline \frac{1}{2}P^T[B^T - A^T H] & \sum_k \lambda_k S_k \end{array} \right] \succeq 0 \right\}$$

is near-optimal:

$$\text{Risk}_{\|\cdot\|}[\hat{x}_{H_*}|\mathcal{X}] \leq \text{Opt} \leq 4 \ln(5[K + L]) \text{RiskOpt}_{\|\cdot\|}[\mathcal{X}],$$

where $\text{RiskOpt}_{\|\cdot\|}[\mathcal{X}] = \inf_{\hat{x}(\cdot)} \text{Risk}_{\|\cdot\|}[\hat{x}|\mathcal{X}]$. is the minimax risk.

Sketch of the proof:

A. The $\|\cdot\|$ -risk of a linear estimate $\hat{x}_H(\omega) = H^T \omega$ is

$$\text{Risk}_{\|\cdot\|}[\hat{x}_H|\mathcal{X}] = \max_{x \in \mathcal{X}} \|[B - H^T A]x\| = \max_{y \in \mathcal{Y}} \|[B - H^T A]Py\|$$

which is nothing but $\|[B - H^T A]P\|_{\|\cdot\|_y \rightarrow \|\cdot\|}$, where $\|\cdot\|_y$ is the norm with the unit ball $\mathcal{Y}$. By the theorem on upper-bounding operator norms, we have

$$\begin{aligned} \text{Risk}_{\|\cdot\|}[\hat{x}_H|\mathcal{X}] &\leq \min_{\lambda, \mu} \left\{ \phi_{\mathcal{T}}(\lambda) + \phi_{\mathcal{R}}(\mu) : \begin{array}{l} \lambda \geq 0, \mu \geq 0 \\ \left[\frac{\sum_{\ell} \mu_{\ell} R_{\ell}}{\frac{1}{2} P^T [B^T - A^T H]} \mid \frac{\frac{1}{2} [B - H^T A] P}{\sum_k \lambda_k S_k} \right] \succeq 0 \end{array} \right\} \\ &\Rightarrow \text{Risk}_{\|\cdot\|}[\hat{x}_{H_*}|\mathcal{X}] \leq \text{Opt}. \end{aligned}$$

B. As we have seen, the quantity

$$\mathfrak{R} = \max_x \{\|Bx\| : Ax = 0, x \in \mathcal{X}\}$$

is a lower bound on the minimal optimal $\|\cdot\|$ -risk $\text{Risk}_{\|\cdot\|}^*[\mathcal{X}]$, and $\mathfrak{R}$ can be tightly (within factor $4 \ln(5[K + L])$) upper-bounded by the optimal value in an explicit conic problem. On the closest inspection, utilizing conic duality, the optimal value in question turns out to be equal to Opt

$\Rightarrow$ The (upper bound on the) $\|\cdot\|$ -risk of the linear estimate $\hat{x}_{H_*}$ is within the factor $4 \ln(5[K + L])$ of the **lower** bound $\mathfrak{R}$ on the minimax optimal $\|\cdot\|$ -risk.