## 10. Multi-objective least squares

- multi-objective least squares
- regularized data fitting
- control
- estimation and inversion


## Multi-objective least squares

we have several objectives

$$
J_{1}=\left\|A_{1} x-b_{1}\right\|^{2}, \quad \ldots, \quad J_{k}=\left\|A_{k} x-b_{k}\right\|^{2}
$$

- $A_{i}$ is an $m_{i} \times n$ matrix, $b_{i}$ is an $m_{i}$-vector
- we seek one $x$ that makes all $k$ objectives small
- usually there is a trade-off: no single $x$ minimizes all objectives simultaneously

Weighted least squares formulation: find $x$ that minimizes

$$
\lambda_{1}\left\|A_{1} x-b_{1}\right\|^{2}+\cdots+\lambda_{k}\left\|A_{k} x-b_{k}\right\|^{2}
$$

- coefficients $\lambda_{1}, \ldots, \lambda_{k}$ are positive weights
- weights $\lambda_{i}$ express relative importance of different objectives
- without loss of generality, we can choose $\lambda_{1}=1$


## Solution of weighted least squares

- weighted least squares is equivalent to a standard least squares problem

$$
\text { minimize }\left\|\left[\begin{array}{c}
\sqrt{\lambda_{1}} A_{1} \\
\sqrt{\lambda_{2}} A_{2} \\
\vdots \\
\sqrt{\lambda_{k}} A_{k}
\end{array}\right] x-\left[\begin{array}{c}
\sqrt{\lambda_{1}} b_{1} \\
\sqrt{\lambda_{2}} b_{2} \\
\vdots \\
\sqrt{\lambda_{k}} b_{k}
\end{array}\right]\right\|^{2}
$$

- solution is unique if the stacked matrix has linearly independent columns
- each matrix $A_{i}$ may have linearly dependent columns (or be a wide matrix)
- it the stacked matrix has linearly independent columns, the solution is

$$
\hat{x}=\left(\lambda_{1} A_{1}^{T} A_{1}+\cdots+\lambda_{k} A_{k}^{T} A_{k}\right)^{-1}\left(\lambda_{1} A_{1}^{T} b_{1}+\cdots+\lambda_{k} A_{k}^{T} b_{k}\right)
$$

## Example with two objectives

$$
\text { minimize }\left\|A_{1} x-b_{1}\right\|^{2}+\lambda\left\|A_{2} x-b_{2}\right\|^{2}
$$

$A_{1}$ and $A_{2}$ are $10 \times 5$

plot shows weighted least squares solution $\hat{x}(\lambda)$ as function of weight $\lambda$

## Example with two objectives

$$
\text { minimize }\left\|A_{1} x-b_{1}\right\|^{2}+\lambda\left\|A_{2} x-b_{2}\right\|^{2}
$$




- left figure shows $J_{1}(\lambda)=\left\|A_{1} \hat{x}(\lambda)-b_{1}\right\|^{2}$ and $J_{2}(\lambda)=\left\|A_{2} \hat{x}(\lambda)-b_{2}\right\|^{2}$
- right figure shows optimal trade-off curve of $J_{2}(\lambda)$ versus $J_{1}(\lambda)$


## Outline

- multi-objective least squares
- regularized data fitting
- control
- estimation and inversion


## Motivation

- consider linear-in-parameters model

$$
\hat{f}(x)=\theta_{1} f_{1}(x)+\cdots+\theta_{p} f_{p}(x)
$$

we assume $f_{1}(x)$ is the constant function 1

- we fit the model $\hat{f}(x)$ to examples $\left(x^{(1)}, y^{(1)}\right), \ldots,\left(x^{(N)}, y^{(N)}\right)$
- large coefficient $\theta_{i}$ makes model more sensitive to changes in $f_{i}(x)$
- keeping $\theta_{2}, \ldots, \theta_{p}$ small helps avoid over-fitting
- this leads to two objectives:

$$
J_{1}(\theta)=\sum_{k=1}^{N}\left(\hat{f}\left(x^{(k)}\right)-y^{(k)}\right)^{2}, \quad J_{2}(\theta)=\sum_{j=2}^{p} \theta_{j}^{2}
$$

primary objective $J_{1}(\theta)$ is sum of squares of prediction errors

## Weighted least squares formulation

$$
\operatorname{minimize} \quad J_{1}(\theta)+\lambda J_{2}(\theta)=\sum_{k=1}^{N}\left(\hat{f}\left(x^{(k)}\right)-y^{(k)}\right)^{2}+\lambda \sum_{j=2}^{p} \theta_{j}^{2}
$$

- $\lambda$ is positive regularization parameter
- equivalent to least squares problem: minimize

$$
\left\|\left[\begin{array}{c}
A_{1} \\
\sqrt{\lambda} A_{2}
\end{array}\right] \theta-\left[\begin{array}{c}
y^{\mathrm{d}} \\
0
\end{array}\right]\right\|^{2}
$$

with $y^{\mathrm{d}}=\left(y^{(1)}, \ldots, y^{(N)}\right)$,

$$
A_{1}=\left[\begin{array}{cccc}
1 & f_{2}\left(x^{(1)}\right) & \cdots & f_{p}\left(x^{(1)}\right) \\
1 & f_{2}\left(x^{(2)}\right) & \cdots & f_{p}\left(x^{(2)}\right) \\
\vdots & \vdots & & \vdots \\
1 & f_{2}\left(x^{(N)}\right) & \cdots & f_{p}\left(x^{(N)}\right)
\end{array}\right], \quad A_{2}=\left[\begin{array}{ccccc}
0 & 1 & 0 & \cdots & 0 \\
0 & 0 & 1 & \cdots & 0 \\
\vdots & \vdots & \vdots & \cdots & \vdots \\
0 & 0 & 0 & \cdots & 1
\end{array}\right]
$$

- stacked matrix has linearly independent columns (for positive $\lambda$ )
- value of $\lambda$ can be chosen by out-of-sample validation or cross-validation


## Example



- solid line is signal used to generate synthetic (simulated) data
- 10 blue points are used as training set; 20 red points are used as test set
- we fit a model with five parameters $\theta_{1}, \ldots, \theta_{5}$ :

$$
\left.\hat{f}(x)=\theta_{1}+\sum_{k=1}^{4} \theta_{k+1} \sin \left(\omega_{k} x+\phi_{k}\right) \quad \text { (with given } \omega_{k}, \phi_{k}\right)
$$

## Result of regularized least squares fit



- minimum test RMS error is for $\lambda$ around 0.08
- increasing $\lambda$ "shrinks" the coefficients $\theta_{2}, \ldots, \theta_{5}$
- dashed lines show coefficients used to generate the data
- for $\lambda$ near 0.08 , estimated coefficients are close to these "true" values


## Outline

- multi-objective least squares
- regularized data fitting
- control
- estimation and inversion


## Control

$$
y=A x+b
$$

- $x$ is $n$-vector of actions or inputs
- $y$ is $m$-vector of results or outputs
- relation between inputs and outputs is a known affine function
the goal is to choose inputs $x$ to optimize different objectives on $x$ and $y$


## Optimal input design

## Linear dynamical system

$$
y(t)=h_{0} u(t)+h_{1} u(t-1)+h_{2} u(t-2)+\cdots+h_{t} u(0)
$$

- output $y(t)$ and input $u(t)$ are scalar
- we assume input $u(t)$ is zero for $t<0$
- coefficients $h_{0}, h_{1}, \ldots$ are the impulse response coefficients
- output is convolution of input with impulse response


## Optimal input design

- optimization variable is the input sequence $x=(u(0), u(1), \ldots, u(N))$
- goal is to track a desired output using a small and slowly varying input


## Input design objectives

$$
\operatorname{minimize} \quad J_{\mathrm{t}}(x)+\lambda_{\mathrm{v}} J_{\mathrm{v}}(x)+\lambda_{\mathrm{m}} J_{\mathrm{m}}(x)
$$

- primary objective: track desired output $y_{\text {des }}$ over an interval $[0, N]$ :

$$
J_{\mathrm{t}}(x)=\sum_{t=0}^{N}\left(y(t)-y_{\mathrm{des}}(t)\right)^{2}
$$

- secondary objectives: use a small and slowly varying input signal:

$$
J_{\mathrm{m}}(x)=\sum_{t=0}^{N} u(t)^{2}, \quad J_{\mathrm{V}}(x)=\sum_{t=0}^{N-1}(u(t+1)-u(t))^{2}
$$

## Tracking error

$$
\begin{aligned}
J_{\mathrm{t}}(x) & =\sum_{t=0}^{N}\left(y(t)-y_{\mathrm{des}}(t)\right)^{2} \\
& =\left\|A_{\mathrm{t}} x-b_{\mathrm{t}}\right\|^{2}
\end{aligned}
$$

with

$$
A_{\mathrm{t}}=\left[\begin{array}{cccccc}
h_{0} & 0 & 0 & \cdots & 0 & 0 \\
h_{1} & h_{0} & 0 & \cdots & 0 & 0 \\
h_{2} & h_{1} & h_{0} & \cdots & 0 & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
h_{N-1} & h_{N-2} & h_{N-3} & \cdots & h_{0} & 0 \\
h_{N} & h_{N-1} & h_{N-2} & \cdots & h_{1} & h_{0}
\end{array}\right], \quad b_{\mathrm{t}}=\left[\begin{array}{c}
y_{\operatorname{des}}(0) \\
y_{\operatorname{des}}(1) \\
y_{\operatorname{des}}(2) \\
\vdots \\
y_{\operatorname{des}}(N-1) \\
y_{\operatorname{des}}(N)
\end{array}\right]
$$

## Input variation and magnitude

Input variation

$$
J_{\mathrm{V}}(x)=\sum_{t=0}^{N-1}(u(t+1)-u(t))^{2}=\|D x\|^{2}
$$

with $D$ the $N \times(N+1)$ matrix

$$
D=\left[\begin{array}{rrrlrrr}
-1 & 1 & 0 & \cdots & 0 & 0 & 0 \\
0 & -1 & 1 & \cdots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & -1 & 1 & 0 \\
0 & 0 & 0 & \cdots & 0 & -1 & 1
\end{array}\right]
$$

Input magnitude

$$
J_{\mathrm{m}}(x)=\sum_{t=0}^{N} u(t)^{2}=\|x\|^{2}
$$

## Example

$\lambda_{\mathrm{v}}=0$,
small $\lambda_{\mathrm{m}}$




## Outline

- multi-objective least squares
- regularized data fitting
- control
- estimation and inversion


## Estimation

Linear measurement model

$$
y=A x_{\mathrm{ex}}+v
$$

- $n$-vector $x_{\text {ex }}$ contains parameters that we want to estimate
- $m$-vector $v$ is unknown measurement error or noise
- $m$-vector $y$ contains measurements
- $m \times n$ matrix $A$ relates measurements and parameters

Least squares estimate: use as estimate of $x_{\text {ex }}$ the solution $\hat{x}$ of

$$
\operatorname{minimize}\|A x-y\|^{2}
$$

## Regularized estimation

add other terms to $\|A x-y\|^{2}$ to include information about parameters

## Example: Tikhonov regularization

$$
\text { minimize }\|A x-y\|^{2}+\lambda\|x\|^{2}
$$

- goal is to make $\|A x-y\|$ small with small $x$
- equivalent to solving

$$
\left(A^{T} A+\lambda I\right) x=A^{T} y
$$

- solution is unique (if $\lambda>0$ ) even when $A$ has linearly dependent columns


## Signal denoising

- observed signal $y$ is $n$-vector

$$
y=x_{\mathrm{ex}}+v
$$

- $x_{\mathrm{ex}}$ is unknown signal
- $v$ is noise


Least squares denoising: find estimate $\hat{x}$ by solving

$$
\text { minimize }\|x-y\|^{2}+\lambda \sum_{i=1}^{n-1}\left(x_{i+1}-x_{i}\right)^{2}
$$

goal is to find slowly varying signal $\hat{x}$, close to observed signal $y$

## Matrix formulation

$$
\text { minimize }\left\|\left[\begin{array}{c}
I \\
\sqrt{\lambda} D
\end{array}\right] x-\left[\begin{array}{l}
y \\
0
\end{array}\right]\right\|^{2}
$$

- $D$ is $(n-1) \times n$ finite difference matrix

$$
D=\left[\begin{array}{rrrlrrr}
-1 & 1 & 0 & \cdots & 0 & 0 & 0 \\
0 & -1 & 1 & \cdots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & & \vdots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & -1 & 1 & 0 \\
0 & 0 & 0 & \cdots & 0 & -1 & 1
\end{array}\right]
$$

- equivalent to linear equation

$$
\left(I+\lambda D^{T} D\right) x=y
$$

## Trade-off

the two objectives $\|\hat{x}(\lambda)-y\|$ and $\|D \hat{x}(\lambda)\|$ for varying $\lambda$



## Three solutions




- $\hat{x}(\lambda) \rightarrow y$ for $\lambda \rightarrow 0$
- $\hat{x}(\lambda) \rightarrow \mathbf{a v g}(y) \mathbf{1}$ for $\lambda \rightarrow \infty$
- $\lambda \approx 10^{2}$ is good compromise


