

Homework 3

EE 263 Stanford University

Summer 2017

- 1. Fitting a model for hourly temperature.** You are given a set of temperature measurements (in degrees C), $y_t \in \mathbb{R}$, $t = 1, \dots, N$, taken hourly over one week (so $N = 168$). An expert says that over this week, an appropriate model for the hourly temperature is a trend (*i.e.*, a linear function of t) plus a diurnal component (*i.e.*, a 24-periodic component):

$$\hat{y}_t = at + p_t,$$

where $a \in \mathbb{R}$ and $p \in \mathbb{R}^N$ satisfies $p_{t+24} = p_t$, for $t = 1, \dots, N - 24$. We can interpret a (which has units of degrees C per hour) as the warming or cooling trend (for $a > 0$ or $a < 0$, respectively) over the week.

- Explain how to find $a \in \mathbb{R}$ and $p \in \mathbb{R}^N$ (which is 24-periodic) that minimize the RMS value of $y - \hat{y}$.
- Carry out the procedure described in part (a) on the data set found in `tempfit_data.m`. Give the value of the trend parameter a that you find. Plot the model \hat{y} and the measured temperatures y on the same plot. (The matlab code to do this is in the data file, but commented out.)
- Temperature prediction.* Use the model found in part (b) to predict the temperature for the next 24-hour period (*i.e.*, from $t = 169$ to $t = 192$). The file `tempdata.m` also contains a 24 long vector `ytom` with tomorrow's temperatures. Plot tomorrow's temperature and your prediction of it, based on the model found in part (b), on the same plot. What is the RMS value of your prediction error for tomorrow's temperatures?

- 2. Curve-smoothing.** We are given a function $F : [0, 1] \rightarrow \mathbb{R}$ (whose graph gives a curve in \mathbb{R}^2). Our goal is to find another function $G : [0, 1] \rightarrow \mathbb{R}$, which is a *smoothed* version of F . We'll judge the smoothed version G of F in two ways:

- Mean-square deviation from F* , defined as

$$D = \int_0^1 (F(t) - G(t))^2 dt.$$

- Mean-square curvature*, defined as

$$C = \int_0^1 G''(t)^2 dt.$$

We want *both* D and C to be small, so we have a problem with two objectives. In general there will be a trade-off between the two objectives. At one extreme, we can choose $G = F$, which makes $D = 0$; at the other extreme, we can choose G to be an affine function (*i.e.*, to have $G''(t) = 0$ for all $t \in [0, 1]$), in which case $C = 0$. The problem is to identify the optimal trade-off curve between C and D , and explain how to find smoothed functions G on the optimal trade-off curve. To reduce the problem to a finite-dimensional one, we will represent the functions F and G (approximately) by vectors $f, g \in \mathbb{R}^n$, where

$$f_i = F(i/n), \quad g_i = G(i/n).$$

You can assume that n is chosen large enough to represent the functions well. Using this representation we will use the following objectives, which approximate the ones defined for the functions above:

- *Mean-square deviation*, defined as

$$d = \frac{1}{n} \sum_{i=1}^n (f_i - g_i)^2.$$

- *Mean-square curvature*, defined as

$$c = \frac{1}{n-2} \sum_{i=2}^{n-1} \left(\frac{g_{i+1} - 2g_i + g_{i-1}}{1/n^2} \right)^2.$$

In our definition of c , note that

$$\frac{g_{i+1} - 2g_i + g_{i-1}}{1/n^2}$$

gives a simple approximation of $G''(i/n)$. You will only work with this approximate version of the problem, *i.e.*, the vectors f and g and the objectives c and d .

- Explain how to find g that minimizes $d + \mu c$, where $\mu \geq 0$ is a parameter that gives the relative weighting of sum-square curvature compared to sum-square deviation. Does your method always work? If there are some assumptions you need to make (say, on rank of some matrix, independence of some vectors, etc.), state them clearly. Explain how to obtain the two extreme cases: $\mu = 0$, which corresponds to minimizing d without regard for c , and also the solution obtained as $\mu \rightarrow \infty$ (*i.e.*, as we put more and more weight on minimizing curvature).
- Get the file `curve_smoothing.m` from the course web site. This file defines a specific vector f that you will use. Find and plot the optimal trade-off curve between d and c . Be sure to identify any critical points (such as, for example, any intersection of the curve with an axis). Plot the optimal g for the two extreme cases $\mu = 0$ and $\mu \rightarrow \infty$, and for three values of μ in between (chosen to show the trade-off nicely). On your plots of g , be sure to include also a plot of f , say with dotted line type, for reference. Submit your matlab code.

3. Hovercraft with limited range. We have a hovercraft moving in the plane with two thrusters, each pointing through the center of mass, exerting forces in the \mathbf{x} and \mathbf{y} directions with 100% efficiency. The hovercraft has mass 1. The discretized equations of motion for the hovercraft are

$$x(t+1) = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} x(t) + \begin{bmatrix} \frac{1}{2} & 0 \\ 1 & 0 \\ 0 & \frac{1}{2} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix}$$

where x_1 and x_2 are the position and velocity in the \mathbf{x} -direction, and x_3 , x_4 are the position and velocity in the \mathbf{y} -direction. Here

$$u(t) = \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix}$$

is the force acting on the hovercraft for time in the interval $[t, t+1)$. Let the position of the vehicle at time t be $q(t) \in \mathbb{R}^2$.

- a) The hovercraft starts at the origin. We'd like to apply thrust to make it move through points p_1, p_2, p_3 at times t_1, t_2, t_3 , where

$$\begin{array}{ccc} p_1 = \begin{bmatrix} 1 \\ -\frac{1}{2} \end{bmatrix} & p_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} & p_3 = \begin{bmatrix} -\frac{3}{2} \\ 0 \end{bmatrix} \\ t_1 = 6 & t_2 = 40 & t_3 = 50 \end{array}$$

We will run the hovercraft on the time interval $[0, 70]$. We'd like to apply a sequence of inputs $u(0), u(1), \dots, u(70)$ to make the hovercraft position pass through the above sequence of points at the specified times.

We would like to find the sequence of inputs that drives the hovercraft through the desired points which has the minimum cost, given by the sum of the squares of the forces:

$$\sum_{t=0}^{70} \|u(t)\|^2$$

To do this, pick A_{hov} and y_{des} to set this problem up as an equivalent minimum-norm problem, where we would like to find the minimum-norm u_{seq} which satisfies

$$A_{\text{hov}} u_{\text{seq}} = y_{\text{des}}$$

where u_{seq} is the sequence of force inputs

$$u_{\text{seq}} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(70) \end{bmatrix}$$

Plot the trajectory of the hovercraft using this input, and the way-points p_1, \dots, p_3 . Also plot the optimal u against time.

- b) Now we would like to compute the trade-off curve between the accuracy with which the mass passes through the waypoints and the norm of the force used. Let our two objective functions be

$$J_1 = \sum_{i=1}^3 \|q(t_i) - p_i\|^2 = \|A_{\text{hov}} u_{\text{seq}} - y_{\text{des}}\|^2$$

and

$$J_2 = \sum_{t=0}^{70} \|u(t)\|^2$$

By minimizing the weighted sum

$$J_1 + \mu J_2$$

for a range of values of μ , plot the trade-off curve of J_1 against J_2 showing the achievable performance. To generate suitable values of μ , you may find the `logspace` command useful in Matlab; you'll need to pick appropriate maximum and minimum values. This above trade-off curve shows how we can trade-off between how accurately the hovercraft passes through the waypoints and how much input energy is used.

- c) For each of the following values of μ

$$\{ 10^{\frac{p}{2}} \mid p = -2, 0, 2, \dots, 10 \}$$

plot the trajectories all on the same plot, together with the waypoints.

- d) Now suppose we are controlling the hovercraft by radio control, and the maximum range possible between the transmitter and receiver is 2 (in whatever units we are using for distance.) Notice that, if we use the minimum-norm input then the hovercraft passes out of range, both when making its first turn and on the final stretch (between times 50 and 70).

We'd like to do something about this, but trading off the input norm as above doesn't do the right thing; if μ is large then the hovercraft stays within range, but misses the waypoints entirely; if μ is small then it comes close to the waypoints, but goes out of range. Notice that this is particularly a problem on the final stretch between times 50 and 70; explain why this is.

- e) One remedy for this problem is to solve a *constrained multiobjective least-squares* problem. We would like to impose the constraint that

$$A_{\text{hov}} u_{\text{seq}} = y_{\text{des}}$$

that is, achieve zero waypoint error $J_1 = 0$. We can attempt to keep the hovercraft in range by trading off the sum of the squares of the *position*

$$J_3 = \sum_{t=0}^{70} \|q(t)\|^2$$

against input cost J_2 subject to this constraint. To do this, we'll solve

$$\begin{aligned} & \text{minimize} && J_3 + \gamma J_2 \\ & \text{subject to} && A_{\text{hov}} u_{\text{seq}} = y_{\text{des}} \end{aligned}$$

First, find the matrix W so that the cost function is given by

$$J_3 + \gamma J_2 = \|Wu_{\text{seq}}\|^2$$

f) Now we have a problem of the form

$$\begin{aligned} & \text{minimize} && \|Wu\|^2 \\ & \text{subject to} && Au = y_{\text{des}} \end{aligned}$$

This is called a *weighted minimum-norm solution*; the only difference from the usual minimum-norm solution to $Au = y_{\text{des}}$ is the presence of the matrix W , and when $W = I$ the optimal u is just given by $u_{\text{opt}} = A^\dagger y_{\text{des}}$. Show that the solution for general W is

$$u_{\text{opt}} = \Sigma^{-1} A^T (A \Sigma^{-1} A^T)^{-1} y_{\text{des}}$$

where $\Sigma = W^T W$. (One way to do this is using Lagrange multipliers.) Use this to solve the remaining parts of this problem.

g) For each of the following values of γ

$$\{ 10^{\frac{p}{2}} \mid p = 0, 2, 4, \dots, 20 \}$$

Plot the trajectories all on the same plot, together with the waypoints. Explain what you see.

h) By trying different values of γ , you should be able to find a trajectory which just keeps the hovercraft within range. Plot the trajectory of the hovercraft; what is the corresponding value of γ ? Is this the smallest-norm input u that just keeps the hovercraft within range, and drives the hovercraft through the waypoints? Explain why, or why not.

i) For a range of values of γ , plot the trade-off curve of J_3 against J_2 showing the achievable performance.

4. The smoothest input that takes the state to zero. We consider the discrete-time linear dynamical system $x(t+1) = Ax(t) + Bu(t)$, with

$$A = \begin{bmatrix} 1.0 & 0.5 & 0.25 \\ 0.25 & 0 & 1.0 \\ 1.0 & -0.5 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 1.0 \\ 0.1 \\ 0.5 \end{bmatrix}, \quad x(0) = \begin{bmatrix} 25 \\ 0 \\ -25 \end{bmatrix}.$$

The goal is to choose an input sequence $u(0), u(1), \dots, u(19)$ that yields $x(20) = 0$. Among the input sequences that yield $x(20) = 0$, we want the one that is *smoothest*, i.e., that minimizes

$$J_{\text{smooth}} = \left(\frac{1}{20} \sum_{t=0}^{19} (u(t) - u(t-1))^2 \right)^{1/2},$$

where we take $u(-1) = 0$ in this formula. Explain how to solve this problem. Plot the smoothest input u_{smooth} , and give the associated value of J_{smooth} .

5. Optimal dynamic purchasing. You are to complete a large order to buy a certain number, B , of shares in some company. You are to do this over T time periods. (Depending on the circumstances, a single time period could be between tens of milliseconds and minutes.) We will let b_t denote the number of shares bought in time period t , for $t = 1, \dots, T$, so we have $b_1 + \dots + b_T = B$. (The quantities B, b_1, \dots, b_T can all be any real number; $b_t < 0$, for example, means we *sold* shares in the period t . We also don't require b_t to be integers.) We let p_t denote the price per share in period t , so the total cost of purchasing the B shares is $C = p_1 b_1 + \dots + p_T b_T$.

The amounts we purchase are large enough to have a noticeable effect on the price of the shares. The prices change according to the following equations:

$$p_1 = \bar{p} + \alpha b_1, \quad p_t = \theta p_{t-1} + (1 - \theta)\bar{p} + \alpha b_t, \quad t = 2, \dots, T.$$

Here \bar{p} is the base price of the shares and α and θ are parameters that determine how purchases affect the prices. The parameter α , which is positive, tells us how much the price goes up in the current period when we buy one share. The parameter θ , which lies between 0 and 1, measures the *memory*: If $\theta = 0$ the share price has no memory, and the purchase made in period t only affects the price in that period; if θ is 0.5 (say), the effect a purchase has on the price decays by a factor of two between periods. If $\theta = 1$, the price has perfect memory and the price change will persist for all future periods.

If purchases didn't increase the price, the cost of purchasing the shares would always be $\bar{p}B$. The difference between the total cost and this cost, $C - \bar{p}B$, is called the *transaction cost*.

Find the purchase quantities b_1, \dots, b_T that minimize the transaction cost $C - \bar{p}B$, for the particular problem instance with

$$B = 10000, \quad T = 10, \quad \bar{p} = 10, \quad \theta = 0.8, \quad \alpha = 0.00015.$$

Give the optimal transaction cost. Also give the transaction cost if all the shares were purchased in the first period, and the transaction cost if the purchases were evenly spread over the periods (*i.e.*, if 1000 shares were purchased in each period). Compare these three quantities.

You must explain your method clearly, using any concepts from this class, such as least-squares, pseudo-inverses, eigenvalues, singular values, etc. If your method requires that some rank or other conditions to hold, say so. You must also check, in your matlab code, that these conditions are satisfied for the given problem instance.

6. Tax policies. In this problem we explore a dynamic model of an economy, including the effects of government taxes and spending, which we assume (for simplicity) takes place at the beginning of each year. Let $x(t) \in \mathbb{R}^n$ represent the pre-tax economic activity at the beginning of year t , across n sectors, with $x(t)_i$ being the pre-tax activity level in sector i . We let $\tilde{x}(t) \in \mathbb{R}^n$ denote the post-tax economic activity, across n sectors, at the beginning of year t . We will assume that all entries of $x(0)$ are positive, which will imply that all entries of $x(t)$ and $\tilde{x}(t)$ are positive, for all $t \geq 0$.

The pre- and post-tax activity levels are related as follows. The government taxes the sector activities at rates given by $r \in \mathbb{R}^n$, with r_i the tax rate for sector i . These rates all satisfy $0 \leq r_i < 1$. The total government revenue is then $R(t) = r^\top x(t)$. This total revenue is then spent in the sectors proportionally, with $s \in \mathbb{R}^n$ giving the spending proportions in the sectors. These spending proportions satisfy $s_i \geq 0$ and $\sum_{i=1}^n s_i = 1$; the spending in sector i

is $s_i R(t)$. The post-tax economic activity in sector i , which accounts for the government taxes and spending, is then given by

$$\tilde{x}(t)_i = x(t)_i - r_i x(t)_i + s_i R(t), \quad i = 1, \dots, n, \quad t = 0, 1, \dots$$

Economic activity propagates from year to year as $x(t+1) = E\tilde{x}(t)$, where $E \in \mathbb{R}^{n \times n}$ is the input-output matrix of the economy. You can assume that all entries of E are positive.

We let $S(t) = \sum_{i=1}^n x(t)_i$ denote the total economic activity in year t , and we let

$$G = \lim_{t \rightarrow \infty} \frac{S(t+1)}{S(t)}$$

denote the (asymptotic) growth rate (assuming it exceeds one) of the economy.

- a) Explain why the growth rate does not depend on $x(0)$ (unless it exactly satisfies a single linear equation, which we rule out as essentially impossible). Express the growth rate G in terms of the problem data r , s , and E , using ideas from the course. You may assume that a matrix that arises in your analysis is diagonalizable and has a single dominant eigenvalue, *i.e.*, an eigenvalue λ_1 that satisfies $|\lambda_1| > |\lambda_i|$ for $i = 2, \dots, n$. (These assumptions aren't actually needed—they're just to simplify the problem for you.)
- b) Consider the problem instance with data

$$E = \begin{bmatrix} 0.3 & 0.4 & 0.1 & 0.6 \\ 0.2 & 0.3 & 0.7 & 0.2 \\ 0.1 & 0.2 & 0.2 & 0.1 \\ 0.4 & 0.2 & 0.3 & 0.2 \end{bmatrix}, \quad r = \begin{bmatrix} 0.45 \\ 0.25 \\ 0.1 \\ 0.1 \end{bmatrix}, \quad s = \begin{bmatrix} 0.15 \\ 0.3 \\ 0.4 \\ 0.15 \end{bmatrix}.$$

Find the growth rate. Now find the growth rate with the tax rate set to zero, *i.e.*, $r = 0$ (in which case s doesn't matter). You are welcome (even, encouraged) to simulate the economic activity to double-check your answer, but we want the value using the expression found in part (a).