This homework is due on Tuesday, July 21, 2020, at 11:59PM. Self-grades are due on Tuesday, July 28, 2020, at 11:59PM.

1 SVD

Note: Solutions to this problem will be released on Wednesday, but you are still required to submit your own original work for this problem.

Find the singular value decomposition of the following matrix (leave all work in exact form, not decimal):

$$A = \begin{bmatrix} 1 & 0 & -\sqrt{3} \\ \sqrt{3} & 0 & 1 \\ 0 & 3 & 0 \end{bmatrix}$$

a) Find the eigenvalues of A^*A and order them from largest to smallest, $\lambda_1 \geq \lambda_2 \geq \lambda_3$.

Solution

$$A^*A = \begin{bmatrix} 1 & \sqrt{3} & 0 \\ 0 & 0 & 3 \\ -\sqrt{3} & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & -\sqrt{3} \\ \sqrt{3} & 0 & 1 \\ 0 & 3 & 0 \end{bmatrix} = \begin{bmatrix} 4 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & 4 \end{bmatrix}$$

$$\lambda_1 = 9, \quad \lambda_2 = 4, \quad \lambda_3 = 4$$

b) Find orthonormal eigenvectors $\vec{v}_i$ of A^*A (all eigenvectors are mutually orthogonal and have unit length).

Solution

$\lambda_1 = 9$:

$$\text{Null}(A^*A - 9I) = \text{Null}\left(\begin{bmatrix} -5 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -5 \end{bmatrix}\right) = \text{span}\left\{\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}\right\}$$

Alternatively:

$$\begin{bmatrix} 4 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & 4 \end{bmatrix} \begin{bmatrix} v_{11} \\ v_{12} \\ v_{13} \end{bmatrix} = \begin{bmatrix} 9v_{11} \\ 9v_{12} \\ 9v_{13} \end{bmatrix} \implies v_{11} = 0, v_{12} = 1, v_{13} = 0$$

$$\vec{v}_1 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$$

Since $\lambda_2 = \lambda_3$, any two mutually orthogonal unit vectors that are also orthogonal to $\vec{v}_1 = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^*$ will work. For example:

$$\vec{v}_2 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \vec{v}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

c) Find the singular values $\sigma_i = \sqrt{\lambda_i}$. Find the $\vec{u}_i$ vectors from:

$$A\vec{v}_i = \sigma_i \vec{u}_i$$

Solution

$$\sigma_1 = 3, \quad \sigma_2 = 2, \quad \sigma_3 = 2$$

$$\vec{u}_1 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad \vec{u}_2 = \begin{bmatrix} \frac{1}{2} \\ \frac{\sqrt{3}}{2} \\ 0 \end{bmatrix}, \quad \vec{u}_3 = \begin{bmatrix} -\frac{\sqrt{3}}{2} \\ \frac{1}{2} \\ 0 \end{bmatrix}$$

d) Write out A as a weighted sum of rank 1 matrices:

$$A = \sigma_1 \vec{u}_1 \vec{v}_1^* + \sigma_2 \vec{u}_2 \vec{v}_2^* + \sigma_3 \vec{u}_3 \vec{v}_3^*$$

Solution

$$A = 3 \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} + 2 \begin{bmatrix} \frac{1}{2} & 0 & 0 \\ \frac{\sqrt{3}}{2} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + 2 \begin{bmatrix} 0 & 0 & -\frac{\sqrt{3}}{2} \\ 0 & 0 & \frac{1}{2} \\ 0 & 0 & 0 \end{bmatrix}$$

2 Eigenfaces

In this problem, we will be be exploring the use of PCA to compress and visualize pictures of human faces. We use the images from the data set Labeled Faces in the Wild. Specifically, we use a set of 13,232 images aligned using deep funneling to ensure that the faces are centered in each photo. Each image is a 100×100 image with the face aligned in the center. To turn the image into a vector, we stack each column of pixels in the image on top of each other, and we normalize each pixel value to be between 0 and 1. Thus, a single image of a face is represented by a 10,000 dimensional vector. A vector this size is a bit challenging to work with directly. We combine the vectors from each image into a single matrix so that we can run PCA. For this problem, we will provide you with the first 1,000 principal components, but you can explore how well the images are compressed with fewer components.

Please remember to download the following dataset available here. Then refer to the IPython notebook eigenfaces.ipynb to answer the following questions.

Note: Solutions to this problem will be released on Wednesday, but you are still required to submit your own original work for this problem.

a) We provide you with a randomly selected subset of 1,000 faces from the training set, the first 1,000 principal components, all 13,232 singular values, and the average of all of the faces. What do we need the average of the faces for?

Solution

We need to zero-center the data by subtracting out the average before running PCA. During the reconstruction, we need to add the average back in.

b) We provide you with a set of faces from the training set and compress them using the first 100 principal components. You can adjust the number of principal components used to do the compression between 1 and 1,000. What changes do you see in the compressed images when you used a small number of components and what changes do you see when you use a large number?

Solution

When fewer principal components are used, the images do not differ much from the average face and do not contain many distinguishinig features. This is to be expected, since very small numbers of components will not account for much of the variation found in faces. When more principal components are used, the images more closely resemble the originals.

c) You can visualize each principal component to see what each dimension “adds” to the high-dimensional image. What visual differences do you see in the first few components compared to the last few components?

Solution

The first few principal components are blury images capturing low frequency data. This low frequency data captures some of the broad variation across faces like lighting. The last few components contain high frequency data where small details vary from face to face.

d) By using PCA on the face images, we obtain orthogonal vectors that point in directions of high variance in the original images. We can use these vectors to transform the data into a lower dimensional space and plot the data points. In the notebook, we provide you with code to plot a subset of 1,000 images using the first two principal comonents. Try plotting other components of the data, and see how the shape of the points change. What difference do you see in the plot when you use the first two principal components compared with the last two principal components? What do you think is the cause of this difference?

Solution

The variance of the points in the plot is larger for the first two components compared to the last two components. We can also confirm that the variance is larger for the first few components because the singular values are large while the singular values for the last few components are small. This happens because PCA orders the principal components by the singular values, which can be used to measure the variability in the data for each component.

e) We can use the principal components to generate new faces randomly. We accomplish this by picking a random point in the low-dimensional space and then multiplying it by the matrix of principal components. In the notebook, we provide you with code to generate faces using the first 1,000 principal components. You can adjust the number of components used. How does this affect the resulting images?

Solution

When fewer components are used, the faces appear more similar and when a very small number of principal components are used, they are almost indistinguishable. When we use more principal components, the synthesized faces appear more distinct. This happens because we are adding more degrees of freedom to our "face" vector when we add more principal components. This allows us to generate faces with more variety because we have more parameters that control how the face looks.

3 Brain-machine interface

The iPython notebook pca_brain_machine_interface.ipynb will guide you through the process of analyzing brain machine interface data using principle component analysis (PCA). This will help you to prepare for the project, where you will need to use PCA as part of a classifier that will allow you to use voice or music inputs to control your car.

Please complete the notebook by following the instructions given.

4 Inverted Pendulum on a Rolling Cart (Mechanical)

Consider the inverted pendulum depicted below, which is placed on a rolling cart and whose equations of motion are given by:

$$\ddot{y} = \frac{1}{\frac{M}{m} + \sin^2\theta}\left(\frac{u}{m} + \dot{\theta}^2\ell\sin\theta - g\sin\theta\cos\theta\right)$$
$$\ddot{\theta} = \frac{1}{\ell(\frac{M}{m} + \sin^2\theta)}\left(-\frac{u}{m}\cos\theta - \dot{\theta}^2\ell\cos\theta\sin\theta + \frac{M+m}{m}g\sin\theta\right).$$

where we use $\dot{x}$ to denote the time derivative of x; that is, $\dot{y} = \frac{dy}{dt}$, $\dot{\theta} = \frac{d\theta}{dt}$, $\ddot{y} = \frac{d^2y}{dt^2}$ and $\ddot{\theta} = \frac{d^2\theta}{dt^2}$.

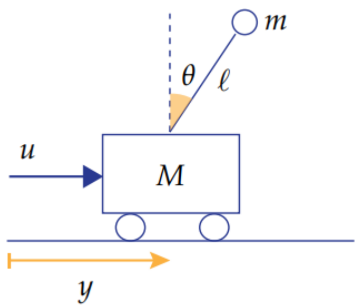

The problems below will prepare us for a future homework problem where we will design a control algorithm to stabilize the upright position.

a) Write the state model using the variables $x_1(t) = \theta(t)$, $x_2(t) = \dot{\theta}(t)$, and $x_3(t) = \dot{y}(t)$. We do not include $y(t)$ as a state variable because we are interested in stabilizing at the point $\theta = 0$, $\dot{\theta} = 0$, $\dot{y} = 0$, and we are not concerned about the final value of the position $y(t)$.

Solution

We have

$$\begin{aligned}
\dot{x_1} &= x_2 && \triangleq f_1(x_1, x_2, x_3, u)\\
\dot{x_2} &= \left(\frac{1}{l(\frac{M}{m}+\sin^2(x_1))}\right)\left(-\frac{u}{m}\cos(x_1) - x_2^2 l\cos(x_1)\sin(x_1) + \frac{M+m}{m}g\sin(x_1)\right) && \triangleq f_2(x_1, x_2, x_3, u)\\
\dot{x_3} &= \left(\frac{1}{\frac{M}{m}+\sin^2(x_1)}\right)\left(\frac{u}{m} + x_2^2 l\sin(x_1) - g\sin(x_1)\cos(x_1)\right) && \triangleq f_3(x_1, x_2, x_3, u)
\end{aligned}$$

b) Show that $x_1 = 0$, $x_2 = 0$, $x_3 = 0$ is an equilibrium point with $u = 0$.

Solution

We check that $\dot{x}_1 = \dot{x}_2 = \dot{x}_3 = 0$ with at the given point:

$$\dot{x}_1 = x_2 = 0$$

$$\dot{x}_2 = \left(\frac{1}{l\left(\frac{M}{m} + \sin^2(0)\right)}\right)\left(-\frac{0}{m}\cos(0) - 0 + \frac{M+m}{m} g \sin(0)\right)$$

$$= \left(\frac{1}{l\left(\frac{M}{m}\right)}\right)(0 - 0 + 0) = 0$$

$$\dot{x}_3 = \left(\frac{1}{\frac{M}{m} + \sin^2(0)}\right)\left(\frac{0}{m} + 0^2 - g\sin(0)\cos(0)\right) = 0.$$

c) Linearize this model at the equilibrium $x_1 = 0$, $x_2 = 0$, $x_3 = 0$, and $u = 0$, and indicate the resulting A and B matrices.

Solution

We can keep in mind that $x_1 = x_2 = x_3 = 0$ to make the derivative much easier. Since we aren't asked to linearize about a particular input, we can linearize about $u^* = 0$. This is fine because f_2 and f_3 are affine (linear plus a constant term) with respect to u.

Note that

$$\begin{array}{lll}
\frac{\partial f_1}{\partial x_1}(0,0,0,0) = 0 & \frac{\partial f_1}{\partial x_2}(0,0,0,0) = 1 & \frac{\partial f_1}{\partial x_3}(0,0,0,0) = 0 \\
\frac{\partial f_2}{\partial x_1}(0,0,0,0) = \frac{M+m}{lM} g & \frac{\partial f_2}{\partial x_2}(0,0,0,0) = 0 & \frac{\partial f_2}{\partial x_3}(0,0,0,0) = 0 \\
\frac{\partial f_3}{\partial x_1}(0,0,0,0) = -\frac{m}{M} g & \frac{\partial f_3}{\partial x_2}(0,0,0,0) = 0 & \frac{\partial f_3}{\partial x_3}(0,0,0,0) = 0,
\end{array}$$

and,

$$\frac{\partial f_1}{\partial u}(0,0,0,0) = 0 \quad \frac{\partial f_2}{\partial u}(0,0,0,0) = -\frac{1}{lM} \quad \frac{\partial f_3}{\partial u}(0,0,0,0) = \frac{1}{M}.$$

Since $x^* = 0$ and $u^* = 0$, we can use the same state variables x and u, instead of $\tilde{x}$ and $\tilde{u}$. Then,

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \approx \underbrace{\begin{bmatrix} 0 & 1 & 0 \\ \frac{M+m}{lM} g & 0 & 0 \\ -\frac{m}{M} g & 0 & 0 \end{bmatrix}}_{A}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ -\frac{1}{lM} \\ \frac{1}{M} \end{bmatrix}}_{B} u.$$

5 Linearizing for understanding amplification

Linearization isn't just something that is important for control, robotics, machine learning, and optimization — it is one of the standard tools for circuits

The circuit below is a voltage amplifier, where the element inside the box is a bipolar junction transistor (BJT).

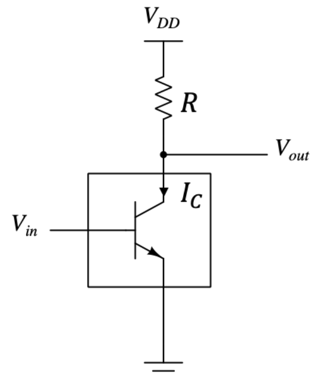

The bipolar transistor in the circuit can be modeled quite accurately as a voltage-controlled current source with

$$I_C(V_{\text{in}}) = I_S e^{\frac{V_{\text{in}}}{V_{\text{TH}}}} \tag{1}$$

where V_{TH} is the thermal voltage. We can assume $V_{TH} = 26$ mV at temperatures of 300K (close to room temperature). I_S is a constant we will try to eliminate.

This BJT can changes small variations in the input voltage V_{in} into large variations in the output voltage V_{out}. We're going to investigate this amplification using linearization.

Let's consider the 2N3904 transistor,[1] where the above expression for $I_C(V_{\text{in}})$ holds as long as $0.2\text{V} < V_{\text{out}} < 40\text{V}$, and $0.1\text{mA} < I_C < 10\text{mA}$.

a) Write a symbolic expression for V_{out} as a function of I_C.

Solution

$V_{\text{out}} = V_{DD} - RI_C$ since we have a voltage drop of $I_C R$ across the resistor and the top voltage is V_{DD}.

b) Now let's linearize I_C in the neighborhood of an input voltage V^*_{in} and a specific I^*_C that satisfy Eq. (1):

$$I_C(V_{in}) = I_C(V^*_{in}) + \delta I_C \approx I^*_C + g_m(V_{in} - V^*_{in}) = I^*_C + g_m\, \delta V_{in} \tag{2}$$

where $\delta V_{in} = V_{in} - V^*_{in}$ is the change in input voltage and $\delta I_C = I_C - I^*_C$ is the change in collector current.

What is g_m here as a function of I^*_C and V_{TH}?[2]

[1] A cheap transistor that people often use in personal projects. You can get them for 3 cents each if you buy in bulk.

[2] If you take EE105, you will learn that this g_m is called the transconductance, and is the single most important parameter in most analog circuit designs.

Solution

We start out by writing out the linearization form that we are looking for.

$$I_C(V_{in}) = I_C^* + \delta I_C = I_C(V_{in}^*) + g_m\ \delta V_{in}$$

Here, we can isolate the δI_C term by subtracting $I_C^* = I_C(V_{in}^*)$ from both sides.

$$\delta I_C = g_m\ \delta V_{in}$$

Now, the meaning of the g_m is the slope of the I_C curve at V_{in}^*.

$$\begin{aligned} g_m &= dI_C/dV_{in}|_{V_{in}^*} && (3) \\ &= \frac{1}{V_{TH}} I_S e^{\frac{V_{in}}{V_{TH}}}|_{V_{in}^*} && (4) \\ &= \frac{I_C^*}{V_{TH}} && (5) \end{aligned}$$

where in the last line, we recognize that the expression in the exponential with the I_S before it is just I_C itself. This is why the I_S constant didn't need to be told to you.

We can use these equations to linearize I_c at certain chosen values of V_{in}, such as values $V_{in}^* = 0.65$ V and $V_{in}^* = 0.65$ V given in parts (d) and (e) below. We plot these linearizations here to help visualize our results.

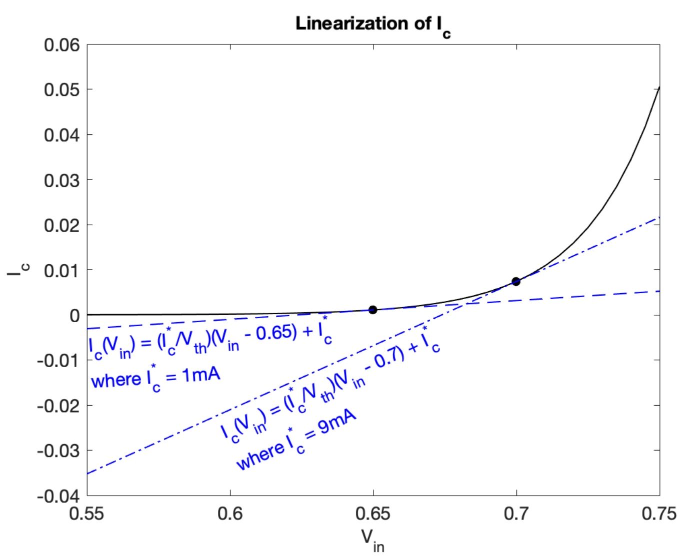

Figure 1: Linearization of I_c

c) We now have a linear relationship, $\delta I_C = g_m\ \delta V_{in}$, between displacements in current and voltage from a known solution (I_C^*, V_{in}^*), called a "bias point" in circuits terminology.

Going back to your equation from part (a), plug in your linearized equation for I_C. Define the appropriate V_{out}^* so that it makes sense to view $V_{out} = V_{out}^* + \delta V_{out}$ when we have $V_{in} = V_{in}^* + \delta V_{in}$, and find the approximate linear relationship between δV_{out} and δV_{in}.[3]

[3]The ratio $\frac{\delta V_{out}}{\delta V_{in}}$ is called the small-signal voltage gain of this amplifier around this bias point.

Solution

Expanding out and remembering the equation for V_{out} from above:

$$V_{out} = V^*_{out} + \delta V_{out} = V_{DD} - R(I^*_C + g_m\, \delta V_{in})$$

Therefore, we define $V^*_{out} = V_{DD} - RI^*_C$ and then

$$\delta V_{out} = -R\, g_m\, \delta V_{in}$$

with g_m as above. Namely

$$\delta V_{out} = -\frac{I^*_C R}{V_{TH}}\, \delta V_{in} = -\frac{V_{DD} - V^*_{out}}{V_{TH}}\, \delta V_{in}$$

.

You don't have to simplify it to this point, but this form is useful because it shows you that the gap between the operating point V_{out^*} to the supply rail V_{DD} matters to understand the small-signal gain. We want as much current as possible to make the gain big, but there is a limit to how big the current can get.

We can use these equations to linearize V_{out} at certain chosen values of V_{in}, such as values $V^*_{in} = 0.65$ V and $V^*_{in} = 0.65$ V given in parts (d) and (e) below. We plot these linearizations here to help visualize our results. The slope of these lines are the small signal voltage gain $\frac{\delta V_{out}}{\delta V_{in}} = -\frac{I^*_C R}{V_{th}}$.

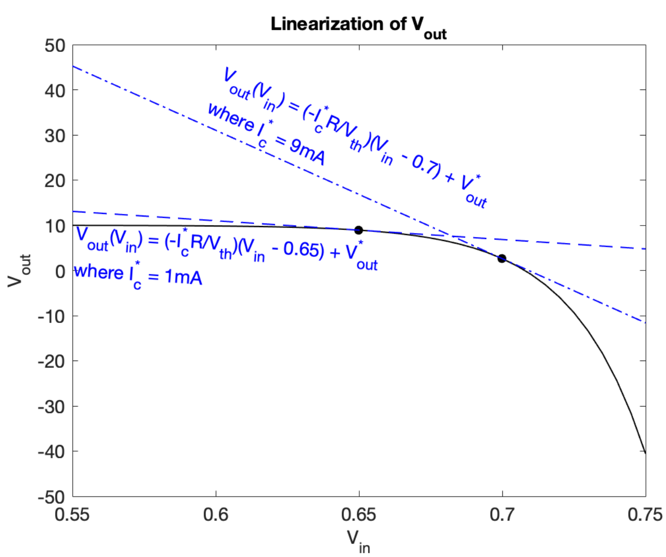

Figure 2: Linearization of V_{out}

6 Homework Process and Study Group

Citing sources and collaborators are an important part of life, including being a student! We also want to understand what resources you find helpful and how much time homework is taking, so we can change things in the future if possible.

a) **What sources (if any) did you use as you worked through the homework?**

b) **If you worked with someone on this homework, who did you work with?** List names and student ID's. (In case of homework party, you can also just describe the group.)

c) **How did you work on this homework?** (For example, *I first worked by myself for 2 hours, but got stuck on problem 3, so I went to office hours. Then I went to homework party for a few hours, where I finished the homework.*)

d) **Do you have any feedback on this homework assignment?**

e) **Roughly how many total hours did you work on this homework?**

f) **How much have you been interacting with your study group? Feel free to describe your interactions.**