Robot Learning
Weekly Exercise 1

Marc Toussaint
Learning & Intelligent Systems Lab, TU Berlin
Marchstr. 23, 10587 Berlin, Germany

Summer 2024

All 4 exercises are a bit too much for a start. Question 3 is bonus.

1 Basic Inverse Kinematics

a) Inverse kinematics (or general constraint solving) can be framed as the optimization problem

$$\min_{q \in \mathbb{R}^n} \|q - q_0\|^2 + \mu \|\phi(q)\|^2 \ , \tag{1}$$

for some constraint function $\phi : \mathbb{R}^n \to \mathbb{R}^d$. Assuming linear $\phi(q) = \phi(q_0) + J(q - q_0)$ with Jacobian J, the solution is

$$q^* = q_0 - (J^\top J + \tfrac{1}{\mu}\mathbf{I})^{-1} J^\top \ \phi(q_0) \ . \tag{2}$$

Verify this by deriving it step by step.

b) To enforce a hard constraint, we want to take the limit $\mu \to \infty$. But $J^\top J$ is typically not invertible (e.g., $d < n$), and you can't directly take the limit in the above solution. However, the solution to this limit is

$$q^* = q_0 - J^\top (J J^\top)^{-1} \phi(q_0) \ . \tag{3}$$

Derive this from the above. Tip: Learn about the Woodbury identity.

a) Derivation...

b) Cheat sheet: https://www.user.tu-berlin.de/mtoussai/notes/gaussians.pdf

Woodbury (for A, B pos.def.): $(A + J^\top B J)^{-1} J^\top B = A^{-1} J^\top (B^{-1} + J A^{-1} J^\top)^{-1}$

Due to the Woodbury identity, the pseudo inverse can be written in two ways (with $W = \mathbf{I}$):

$$J^\# = (W/\mu + J^\top J)^{-1} J^\top = W^{-1} J^\top (J W^{-1} J^\top + \mathbf{I}/\mu)^{-1} \tag{4}$$

Note that you CANNOT USE THE FIRST VERSION to take the limit $\mu \to \infty$ because $J^\top J$ is not invertible. (It is a $n \times n$-matrix of rank d.) But you can use the second version to let $\mu \to \infty$ and $J^\# \to W^{-1} J^\top (J W^{-1} J^\top)^{-1}$, where $J W^{-1} J^\top$ is a $d \times d$-matrix with full rank (for non-singular J).

2 Point mass under PD control

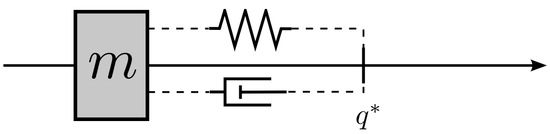

Consider a point mass in a 1D space together with a PD control law:

- The point has mass m, and position $q(t) \in \mathbb{R}$.

- The PD controller applies linear force
$$u(t) = -k_p q(t) - k_d \dot q(t)$$
to the point, where $k_p, k_d \in \mathbb{R}$ are positive constants.

- The resulting dynamics is $m\ddot q(t) = u(t)$.

a) Given the initial state $q(0) = a, \dot q(0) = 0$, what is $q(t)$? (Solve the differential equation.)

Ansatz: Assume $q(t) = c\, e^{\lambda t}$ (where $c, \lambda \in \mathbb{C}$!!)

Let's first solve the differential equation, then later care about boundary constraints $q(0) = a, \dot q(0) = 0$:

$$m\ c\ \lambda^2\ e^{\lambda t} = -k_p\ c\ e^{\lambda t} - k_d\ c\ \lambda e^{\lambda t} \tag{5}$$
$$0 = \left[m\ c\ \lambda^2 + k_d\ c\ \lambda\ + k_p\ c\right] e^{\lambda t} \tag{6}$$
$$0 = m\ \lambda^2 + k_d\ \lambda + k_p \tag{7}$$
$$\lambda = \frac{-k_d \pm \sqrt{k_d^2 - 4mk_p}}{2m} \tag{8}$$

The term $-\frac{k_d}{2m}$ in λ is real $\leftrightarrow$ exponential decay

The square root is (typically) negative $\leftrightarrow$ oscilatory, with $\pm$ just orientation

(I DIDN'T LOOK AT THE OVERDAMPED CASE.)

Now let's look at the boundary conditions: Let's write $c = a + i\bar a$ with $a, \bar a \in \mathbb{R}$:

$$a = q(t) = Re(c) = a \tag{9}$$
$$0 = \dot q(t) = Re(c\lambda) = \frac{-ak_d \pm \bar a\sqrt{|k_d^2 - 4mk_p|}}{2m} \tag{10}$$
$$\bar a = \pm\frac{ak_d}{\sqrt{|k_d^2 - 4mk_p|}} \tag{11}$$

and the velocity constraint can be realized just by a phase shift by $\bar a$.

(I think I now get where you got the idea of "overlaying sin and cos solutions" from... In the complex notation, that's just a phase shift.)

b) The solution describes a damped oscillation around the set-point $q^* = 0$. How do you have to choose k_p and k_d such that the behavior becomes the exponential approach $q(t) = ae^{-t/\tau}$ for some time scale $\tau \in \mathbb{R}$? (This is called "critically damped".)

$$k_p = m/\tau^2\ , k_d = 2m\xi/\tau$$

In general $\xi \in (0, 1]$ gives the damping coefficient. $\xi = 1$ is critically damped

3 BONUS: Fun with Euler-Lagrange

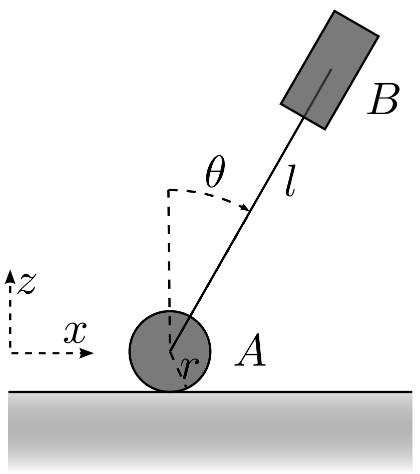

Consider an inverted pendulum mounted on a wheel in the 2D x-z-plane; similar to a Segway. The exercise is to derive the Euler-Lagrange equation for this system.

a) Describe the **pose** $p_i \in \mathbb{R}^3$ of every body in (x, z, ϕ) coordinates: its position in the x-z-plane, and its rotation ϕ relative to the world-vertical. Assume fixed parameters r: radius of the wheel, l: length of the pendulum (height of its COM).

$$p_A = \begin{pmatrix} x \\ 0 \\ \frac{x}{r} \end{pmatrix}, \qquad p_B = \begin{pmatrix} x + \sin(\theta)l \\ \cos(\theta)l \\ \theta \end{pmatrix} \tag{12}$$

b) Describe the (linear and angular) velocity $v_i = \dot{p}_i \in \mathbb{R}^3$ of every body.

$$v_A = \begin{pmatrix} \dot{x} \\ 0 \\ \frac{\dot{x}}{r} \end{pmatrix}, \qquad v_B = \begin{pmatrix} \dot{x} + \dot{\theta}\cos(\theta)l \\ -\dot{\theta}\sin(\theta)l \\ \dot{\theta} \end{pmatrix} \tag{13}$$

c) Formulate the total kinetic energy $T = \frac{1}{2}\sum_i v_i^\top M_i v_i$, summing over the two bodies $i = A, B$. Note that

$$M_i = \begin{pmatrix} m_i & 0 & 0 \\ 0 & m_i & 0 \\ 0 & 0 & I_i \end{pmatrix} \tag{14}$$

with $m_i \in \mathbb{R}$ the normal mass of body i, and $I_i \in \mathbb{R}$ the rotational inertia of body i.

$$T = \frac{1}{2}v_A^\top M_A v_A + \frac{1}{2}v_B^\top M_B v_B \tag{15}$$

$$= \frac{1}{2}\left(\dot{x}^2 m_A + \frac{\dot{x}^2}{r^2}I_A + \dot{x}^2 m_B + 2m_B\dot{x}\dot{\theta}\cos(\theta)l + m_B\dot{\theta}^2\cos(\theta)^2 l^2 + m_B\sin(\theta)^2\dot{\theta}^2 l^2 + \dot{\theta}^2 I_B\right) \tag{16}$$

$$= \frac{1}{2}\left(\dot{x}^2(m_A + m_B + \frac{I_A}{r^2}) + 2m_B\dot{x}\dot{\theta}\cos(\theta)l + \dot{\theta}^2(m_B l^2 + I_B)\right) \tag{17}$$

d) Formulate the potential energy U

$$U = gm_B\cos(\theta)l \tag{18}$$

e) Bonus: Compute the Euler-Lagrange Equation

$$u = \frac{d}{dt}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q}, \tag{19}$$

with $L = T - U$, using the minimal coordinates $q = (x, \theta)$, where x is the position of the wheel and θ the angle of the pendulum relative to the world-vertical.

$$L = T - U \tag{20}$$

$$= \frac{1}{2}\left(\dot{x}^2(m_A + m_B + \frac{I_A}{r^2}) + 2m_B\dot{x}\dot{\theta}\cos(\theta)l + \dot{\theta}^2(m_B l^2 + I_B)\right) - gm_B\cos(\theta)l \tag{21}$$

$$\frac{\partial L}{\partial x} = 0 \tag{22}$$

$$\frac{\partial L}{\partial \dot{x}} = \dot{x}(m_A + m_B + \frac{I_A}{r^2}) + m_B\dot{\theta}\cos(\theta)l \tag{23}$$

$$\frac{\partial L}{\partial \theta} = -m_B\dot{x}\dot{\theta}\sin(\theta)l + gm_B\sin(\theta)l = m_B l\sin(\theta)(g - \dot{x}\dot{\theta}) \tag{24}$$

$$\frac{\partial L}{\partial \dot{\theta}} = \dot{\theta}(m_B l^2 + I_B) + m_B\dot{x}\cos(\theta)l \tag{25}$$

$$\tau_1 = \frac{d}{dt}\frac{\partial L}{\partial \dot{x}} - \frac{\partial L}{\partial x} = \ddot{x}(m_A + m_B + \frac{I_A}{r^2}) + m_B\ddot{\theta}\cos(\theta)l - m_B\dot{\theta}^2\sin(\theta)l \tag{26}$$

$$\tau_2 = \frac{d}{dt}\frac{\partial L}{\partial \dot\theta} - \frac{\partial L}{\partial \theta} = \ddot\theta(m_B l^2 + I_B) + m_B \ddot x \cos(\theta) l - m_B \dot x \dot\theta \sin(\theta) l - m_B l \sin(\theta)(g - \dot x \dot\theta) \tag{27}$$

$$= \ddot\theta(m_B l^2 + I_B) + m_B \ddot x \cos(\theta) l - m_B l \sin(\theta) g \tag{28}$$

4 Logistic Regression

Consider a binary classification problem with data $D = \{(x_i, y_i)\}_{i=1}^n$, $x_i \in \mathbb{R}^d$ and $y_i \in \{0, 1\}$. We define

$$f(x) = x^\top \beta \tag{29}$$

$$p(x) = \sigma(f(x)) \ , \quad \sigma(z) = 1/(1 + e^{-z}) \tag{30}$$

$$L^{\text{nll}}(\beta) = -\sum_{i=1}^n \Big[y_i \log p(x_i) + (1 - y_i) \log[1 - p(x_i)] \Big] \tag{31}$$

where $\beta \in \mathbb{R}^d$ is the model parameter, $\sigma(z)$ the sigmoidal function, and $L^{\text{nll}}(\beta)$ the neg-log-likelihood of the data under the model.

a) Compute the derivative $\frac{\partial}{\partial \beta} L(\beta)$. Tip: use the fact $\frac{\partial}{\partial z}\sigma(z) = \sigma(z)(1 - \sigma(z))$.

b) Compute the 2nd derivative $\frac{\partial^2}{\partial \beta^2} L(\beta)$.

c) How is the neg-log-likelihood related to the cross-entropy? How would the above change when adding an additional regularization $\lambda \|\beta\|^2$ to the loss?

$$L(\beta) = -\sum_{i=1}^n \log P(y_i \,|\, x_i) + \lambda \|\beta\|^2 \tag{32}$$

$$= -\sum_{i=1}^n \Big[y_i \log p_i + (1 - y_i) \log[1 - p_i] \Big] + \lambda \|\beta\|^2 \tag{33}$$

$$\nabla L(\beta) = \frac{\partial L(\beta)}{\partial \beta}^\top = \sum_{i=1}^n (p_i - y_i)\ x_i + 2\lambda I \beta = X^\top (p - y) + 2\lambda I \beta \tag{34}$$

$$\nabla^2 L(\beta) = \frac{\partial^2 L(\beta)}{\partial \beta^2} = \sum_{i=1}^n p_i (1 - p_i)\ x_i\ x_i^\top + 2\lambda I = X^\top W X + 2\lambda I \tag{35}$$

$$\text{where } p(x) := P(y = 1 \,|\, x) = \sigma(x^\top \beta),\ p_i := p(x_i),\ W := \text{diag}(p \circ (1 - p)) \tag{36}$$

(iii) same! nnl=cross-entropy with one-hot encoded target; (above includes λ)