

Exercise 10: Model Predictive Control with `acados`

Prof. Dr. Moritz Diehl and teaching assistant(s)

In this exercise, we will implement a model predictive control (MPC) based controller, which repeatedly solves optimal control problems (OCP) as part of a feedback loop. The controller will be implemented using `acados`, and as control task we consider the swing-up of a pendulum mounted on top of a cart. We will then further speed up the controller sampling time by using the real time iteration scheme (RTI).

`acados` is a software package providing fast embedded solvers for nonlinear optimal control. Problems are formulated based on CasADi symbolics and the `acados` OCP interface. Among others, `acados` provides numerical integrators for ordinary and differential algebraic equations (ODE resp. DAE), based on both implicit and explicit Runge-Kutta methods, as well as sequential quadratic programming (SQP)-type OCP solvers. The `acados` core library is implemented in C, but can be accessed via interfaces to Python, Octave, and Matlab. Given an OCP description, `acados` generates self-contained C code, which – after compilation – can be executed efficiently to solve the problem. For more information, see <https://docs.acados.org/>.

I Installing `acados`

Install `acados` by following the instruction at <https://docs.acados.org/installation/index.html> and the Python interface by following <https://docs.acados.org/interfaces/index.html>. To ensure everything works properly, execute `examples/acados_python/getting_started/minimal_example_ocp.py` from the examples folder. This should run without error.

As control task, we consider the swing-up of a pendulum mounted on a cart, as illustrated below:

The state consists of the horizontal position of the cart center p_x , the angle of the pendulum θ , as well as the respective velocities v_x and ω . The pendulum has length l with a mass m mounted at its tip. The cart has mass M and can be actuated by applying a horizontal force F , i.e., $u = F$. The corresponding dynamics are

$$x = \begin{pmatrix} p_x \\ \theta \\ v_x \\ \omega \end{pmatrix}, \quad \dot{x} = \begin{pmatrix} v_x \\ \omega \\ \frac{-ml \sin \theta \omega^2 + mg \cos \theta \sin \theta + F}{M + m(1 - \cos^2 \theta)} \\ \frac{-ml \cos \theta \sin \theta \omega^2 + F \cos \theta + (M + m)g \sin \theta}{l(M + m(1 - \cos^2 \theta))} \end{pmatrix} := f(x, u). \quad (1)$$

where g is the gravitational acceleration.

By penalizing deviations of state and control from the origin, the swing-up task can be expressed by the continuous-time OCP

$$\min_{x(\cdot), u(\cdot)} \int_0^T \frac{1}{2} x(t)^\top Q x(t) + \frac{1}{2} u(t)^\top R u(t) dt + \frac{1}{2} x(T)^\top Q x(T) \quad (2a)$$

$$\text{s.t.} \quad x(0) = \bar{x}_0, \quad (2b)$$

$$\dot{x}(t) = f(x(t), u(t)), \quad t \in [0, T], \quad (2c)$$

$$-F_{\max} \leq u(t) \leq F_{\max}, \quad t \in [0, T], \quad (2d)$$

with cost matrices $Q \succeq 0$, $R \succeq 0$.

II Formulating and solving the OCP with acados

1. Complete the templates `cartpole_model.py` and `cartpole_ocp.py` in order to solve OCP (2) with `acados`. A detailed overview of the `acados` OCP interface can be found at https://github.com/acados/acados/blob/master/docs/problem_formulation/problem_formulation_ocp_mex.pdf. Specifically, note that:
 - (a) In general, `acados` expects dynamics defined in implicit form via $f_{\text{impl}}(x, \dot{x}, u) = 0$. Thus, in our case we have $f_{\text{impl}}(x, \dot{x}, u) = \dot{x} - f(x, u)$.
 - (b) We use a stage cost of the form 'NONLINEAR_LS' (nonlinear least squares), i.e., $l(x, u) = \frac{1}{2} \|y - y_{\text{ref}}\|_W$, with $y = F(x, u)$ a (possibly) nonlinear expression depending on x and u . In our case, $y = (x, u)$ is simply the concatenation of state and controls, and $W = \text{blockdiag}(Q, R)$, i.e., the blockdiagonal concatenation of Q and R . Similar considerations hold for the terminal cost.
2. Take a detailed look at the above templates files. On a high level, describe the steps `acados` takes to discretize and solve the OCP. Refer to the concepts learned in the course.

III Model predictive control

Since we can now successfully solve the OCP, we are able to construct an MPC loop around the OCP. For this purpose, we read the first discretized control value u_0 from the OCP solution and apply it to the simulation of the controlled system (when deploying the controller in the real world, this would be the actual physical system). After one simulation step, this results in a new state. We resolve the OCP for this new state, and again employ the resulting u_0 , and so on.

1. Complete and run the template `cartpole_closed_loop.py`. How long does it take for the MPC controller to return a control input after receiving the current state?
2. The provided template uses the exact same dynamics for the simulation of the ‘real’ system as are used in the OCP. This is clearly unrealistic, since in reality our model will never be perfect. Create some model-plant mismatch by perturbing the simulation. For this purpose, add some noise at every time step of the simulation, e.g., $w \sim \mathcal{N}(0, \Sigma)$ with $\Sigma = \text{diag}(\sigma_p^2, \sigma_\theta^2, \sigma_v^2, \sigma_\omega^2)$ and $\sigma_p = \sigma_\theta = 10^{-3}$ and $\sigma_v = \sigma_\omega = 10^{-2}$. You can also create some bias by using a nonzero mean. Alternatively, or additionally, you can add strong perturbations at specific time points, corresponding to someone suddenly kicking the pendulum tip. For this purpose, after half the simulation horizon has passed, add $\pi/10$ to the state θ . Play around with the disturbances and see how the controller responds.

*Hint: If you perturb the system too strongly, **acados** will not be able to solve the OCP in the given number of maximum iterations. This results in a warning referring to ‘status 2’ (generated by checking the OCP solver return status). The returned control u_0 will then correspond to a nonconverged solution.*

IV Real-time iteration scheme

Due to model-plant mismatch, it is in general desirable to control systems at a high frequency, in order to be able to quickly react to deviations between prediction and reality. Considering computational delay, there is a trade-off between taking time to fully solve the OCP or quickly returning a non-converged, suboptimal, solution. The main ideas of the real-time iteration scheme (RTI) can be summarized as:

- After receiving the current system state $\bar{x}_0$, only perform a single SQP iteration before returning u_0 .
 - While waiting for the updated system state $\bar{x}_0$, already perform as many computations as possible without knowing $\bar{x}_0$ (“preparation phase”). Then, only few computations need to be performed after receiving $\bar{x}_0$ (“feedback phase”).
1. Consider a discrete-time OCP in simultaneous formulation (multiple shooting), solved with an SQP method. Assume we are currently at iteration i , with current solution guess y^i , where y collects all optimization variables. Which computations can already be performed while the initial state $\bar{x}_0$ is still unspecified?
 2. Run the template `cartpole_closed_loop.py` with RTI turned on. How long does it take for the RTI controller to return a control input after receiving the current state? How does this compare to the fully converged MPC from the previous task? How do the resulting trajectories compare? Why is the latter comparison not fully fair?