Augmented Lagrangian Iterative LQR

Thursday, 5 February 2026 17:59

Numerical Optimizations

Recall the formulation and results that we had for iLQR, except now we are replacing the overall cost function with the augmented version:

$$L_A = l_N(x_N) + \lambda_N c_N(x_N) + \frac{1}{2} c_N(x_N)^\top I_{\mu,N} c_N(x_N) + \sum_{k=0}^{N-1} \left[l_k(x_k, u_k) + \lambda_k c_k(x_k, u_k) + \frac{1}{2} c_k(x_k, u_k)^\top I_{\mu,k} c_k(x_k, u_k) \right]$$

Observation: Constraints at final step do not depend on u

The cost-to-go and action value functions are now:

$$V_N(x_N)\Big|_{\lambda,\mu} = L_N(x_N, \lambda_N, \mu_N)$$

$$V_k(x_k)\Big|_{\lambda,\mu} = \min_{u_k} \left\{ L_k(x_k, u_k, \lambda_k, \mu_k) + V_{k+1}\big(f(x_k, u_k, \Delta t)\big)\Big|_{\lambda,\mu} \right\} = \min_{u_k} Q(x_k, u_k)\Big|_{\lambda,\mu}$$

Let us, again like in iLQR, make $V_k(x_k)$ a quadratic:

$$\delta V_k(x_k) \approx \frac{1}{2} \delta x_k^\top P_k \delta x_k + p_k^\top \delta x_k$$

Terminal Gradient p_N

$$p_N = \frac{\partial V_N}{\partial x}$$

$$p_N = \frac{\partial \bar{l}_f(x_N)}{\partial x}$$

$$p_N = \frac{\partial l_f(x_N)}{\partial x} + \frac{\partial c(x_N)}{\partial x}^\top \lambda + \frac{\partial c(x_N)}{\partial x}^\top I_{\mu,N} c(x_N)$$

Terminal Hessian P_N

Consider the terminal Hessian P_N, defined as

$$P_N = \frac{\partial^2 V_N}{\partial x^2}$$

$$P_N = \frac{\partial^2 \bar{l}_f(x_N)}{\partial x^2}$$

$$P_N = \frac{\partial^2 l_f(x_N)}{\partial x^2} + \frac{\partial c(x_N)}{\partial x}^\top I_{\mu,N} \frac{\partial c(x_N)}{\partial x}$$

Action-Value Taylor Expansion

$$Q(x_k, u_k) = L_k(x_k, u_k, \lambda_k, \mu_k) + V_{k+1}\big(f(x_k, u_k, \Delta t)\big)\Big|_{\lambda,\mu}$$

$l_k(x_k, u_k)$ Taylor Expansion

This remains virtually unchanged from iLQR.

$$\delta l_k(x_k, u_k) \approx \frac{\partial l}{\partial x_k} \delta x_k + \frac{\partial l}{\partial u_k} \delta u_k + \frac{1}{2} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}^\top \begin{bmatrix} \frac{\partial^2 l}{\partial x_k^2} & \frac{\partial^2 l}{\partial x_k \partial u_k} \\ \frac{\partial^2 l}{\partial u_k \partial x_k} & \frac{\partial^2 l}{\partial u_k^2} \end{bmatrix} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}$$

$\lambda_k c_k(x_k, u_k) + \frac{1}{2} c_k(x_k, u_k)^\top I_{\mu,k} c_k(x_k, u_k)$ Taylor Expansion

This is a brand new term that we get attached to the $l_k(x_k, u_k)$.

$$\delta(\dots) \approx \frac{\partial(\dots)}{\partial x_k} \delta x_k + \frac{\partial(\dots)}{\partial u_k} \delta u_k + \frac{1}{2} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}^\top \begin{bmatrix} \frac{\partial^2(\dots)}{\partial x_k^2} & \frac{\partial^2(\dots)}{\partial x_k \partial u_k} \\ \frac{\partial^2(\dots)}{\partial u_k \partial x_k} & \frac{\partial^2(\dots)}{\partial u_k^2} \end{bmatrix} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}$$

$$\delta(\dots) \approx \lambda_k \frac{\partial c_k}{\partial x_k} \delta x_k + \frac{\partial c_k}{\partial x_k} I_{\mu,k} c_k \delta x_k + \lambda_k \frac{\partial c_k}{\partial u_k} \delta x_k + \frac{\partial c_k}{\partial u} I_{\mu,k} c_k \delta u_k + \frac{1}{2} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}^\top \begin{bmatrix} \frac{\partial c_k}{\partial x_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial x_k} & \frac{\partial c_k}{\partial x_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial u_k} \\ \frac{\partial c_k}{\partial u_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial x_k} & \frac{\partial c_k}{\partial u_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial u_k} \end{bmatrix} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}$$

$V_{k+1}(f(x_k, u_k, \Delta t))$ Taylor Expansion

This remains virtually unchanged from iLQR, with the exception that the additional AL terms are now included in the cost function.

$$\delta V_{k+1}\big(f(x_k, u_k)\big) \approx \frac{\partial V_{k+1}}{\partial x}\Big|_{x_{k+1}}^\top \delta x_{k+1} + \frac{1}{2} \delta x_{k+1}^\top \frac{\partial^2 V_{k+1}}{\partial x^2}\Big|_{x_{k+1}} \delta x_{k+1}$$

$$\delta V_{k+1}\big(f(x_k, u_k)\big) \approx p_{k+1} \delta x_{k+1} + \frac{1}{2} \delta x_{k+1}^\top P_{k+1} \delta x_{k+1}$$

$$\delta V_{k+1}\big(f(x_k, u_k)\big) \approx p_{k+1} [A_k \delta x_k + B_k \delta u_k] + \frac{1}{2} [A_k \delta x_k + B_k \delta u_k]^\top P_{k+1} [A_k \delta x_k + B_k \delta u_k]$$

$$\delta V_{k+1}\big(f(x_k, u_k)\big) \approx \begin{bmatrix} A_k^\top p_{k+1} \\ B_k^\top p_{k+1} \end{bmatrix}^\top \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}^\top \begin{bmatrix} A_k^\top P_{k+1} A_k & A_k^\top P_{k+1} B_k \\ B_k^\top P_{k+1} A_k & B_k^\top P_{k+1} B_k \end{bmatrix} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}$$

Combining all Results

$$\delta Q_k = \frac{1}{2} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}^\top \begin{bmatrix} \frac{\partial^2 l}{\partial x_k^2} + A_k^\top P_{k+1} A_k + \frac{\partial c_k}{\partial x_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial x_k} & \frac{\partial^2 l}{\partial x_k \partial u_k} + A_k^\top P_{k+1} B_k + \frac{\partial c_k}{\partial x_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial u_k} \\ \frac{\partial^2 l}{\partial u_k \partial x_k} + B_k^\top P_{k+1} A_k + \frac{\partial c_k}{\partial u_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial x_k} & \frac{\partial^2 l}{\partial u_k^2} + B_k^\top P_{k+1} B_k + \frac{\partial c_k}{\partial u_k}^\top I_{\mu,k} \frac{\partial c_k}{\partial u_k} \end{bmatrix} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix} + \begin{bmatrix} \frac{\partial l}{\partial x_k} + A_k^\top p_{k+1} + \lambda_k \frac{\partial c_k}{\partial x_k} + \frac{\partial c_k}{\partial x_k} I_{\mu,k} c_k \\ \frac{\partial l}{\partial u_k} + B_k^\top p_{k+1} + \lambda_k \frac{\partial c_k}{\partial u_k} + \frac{\partial c_k}{\partial u} I_{\mu,k} c_k \end{bmatrix}^\top \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}$$

$$\delta Q_k = \frac{1}{2} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}^\top \begin{bmatrix} Q_{xx} & Q_{xu} \\ Q_{ux} & Q_{uu} \end{bmatrix} \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix} + \begin{bmatrix} Q_x \\ Q_u \end{bmatrix}^\top \begin{bmatrix} \delta x_k \\ \delta u_k \end{bmatrix}$$

Optimal Adjustment Control

$$\frac{\partial \delta Q_k}{\partial \delta u_k} = Q_u + Q_{ux} \delta x + Q_{uu} \delta u = 0$$

$$\delta u_k = -Q_{uu}^{-1} (Q_{ux} \delta x_k + Q_u)$$

Adding regularization and rewriting as feedback law:

$$\delta u_k = -(Q_{uu} + \rho I)^{-1} (Q_{ux} \delta x_k + Q_u)$$

$$\delta u_k = K_k \delta x_k + d_k$$

Closed-Form Expression

Finally, inserting the control law into δQ_k yields:

$$\delta Q_k = \frac{1}{2} \begin{bmatrix} \delta x_k \\ K_k \delta x_k + d_k \end{bmatrix}^\top \begin{bmatrix} Q_{xx} & Q_{xu} \\ Q_{ux} & Q_{uu} \end{bmatrix} \begin{bmatrix} \delta x_k \\ K_k \delta x_k + d_k \end{bmatrix} + \begin{bmatrix} Q_x \\ Q_u \end{bmatrix}^\top \begin{bmatrix} \delta x_k \\ K_k \delta x_k + d_k \end{bmatrix}$$

$$\delta Q_k = \frac{1}{2} \left[\delta x_k^\top Q_{xx} \delta x_k + \delta x_k^\top Q_{xu} (K_k \delta x_k + d_k) + (K_k \delta x_k + d_k)^\top Q_{ux} \delta x_k + (K_k \delta x_k + d_k)^\top Q_{uu} (K_k \delta x_k + d_k) \right] + Q_x^\top \delta x_k + Q_u^\top (K_k \delta x_k + d_k)$$

$$\delta Q_k = Q_u^\top d_k + \frac{1}{2} d_k^\top Q_{uu} d_k + Q_x^\top \delta x_k + Q_u^\top (K_k \delta x_k) + \frac{1}{2} \left[\delta x_k^\top Q_{xu} d_k + d_k^\top Q_{ux} \delta x_k + (K_k \delta x_k)^\top Q_{uu} d_k + d_k^\top Q_{uu} (K_k \delta x_k) \right] + \frac{1}{2} \left[\delta x_k^\top Q_{xx} \delta x_k + \delta x_k^\top Q_{xu} (K_k \delta x_k) + (K_k \delta x_k)^\top Q_{ux} \delta x_k + (K_k \delta x_k)^\top Q_{uu} (K_k \delta x_k) \right]$$

Two astute observations. First of all, since we solved for $\delta V_k = \min_{\delta u_k} \delta Q_k$, we effectively have:

$$\delta V_k = Q_u^\top d_k + \frac{1}{2} d_k^\top Q_{uu} d_k + Q_x^\top \delta x_k + Q_u^\top (K_k \delta x_k) + d_k^\top Q_{uu} K_k \delta x_k + d_k^\top Q_{ux} \delta x_k + \frac{1}{2} \left[\delta x_k^\top Q_{xx} \delta x_k + \delta x_k^\top Q_{xu} (K_k \delta x_k) + (K_k \delta x_k)^\top Q_{ux} \delta x_k + (K_k \delta x_k)^\top Q_{uu} (K_k \delta x_k) \right]$$

Finally we have:

$$\delta V_k = \Delta V_k + p_k^\top \delta x_k + \frac{1}{2} \delta x_k^\top P_k \delta x_k$$

$$\Delta V_k = Q_u^\top d_k + \frac{1}{2} d_k^\top Q_{uu} d_k$$

$$p_k = Q_x + K_k^\top Q_{uu} d_k + K_k^\top Q_u + Q_{xu} d_k$$

$$P_k = Q_{xx} + K_k^\top Q_{uu} K_k + K_k^\top Q_{ux} + Q_{xu} K_k$$

Basically, we are using Bellman's Optimality Principle to go backwards from the final step to the first step. Ideally, we want to find a closed-form solution for this.

We start by expressing the cost at timestep k+1 as a quadratic. Next, we want to find the action-value at k, which depends on the cost at k+1. We do this by making everything a second-order quadratic, and removing the dependence on u_k by setting the first derivative to 0.

As it turns out, this does two things:

- It fulfills cost = min_u action-value, meaning that our result with inserted feedback law is now the cost and not just the action-value
- The resulting cost is also quadratic. It has a constant term, but for optimization, any derivative removes this, so the results carry over