MACHINE LEARNING

WEEK 3: OPTIMIZATION

LECTURE 2: HAMILTONIAN, KALMAN FILTER

Euler-Lagrange Method

$$\max_{\boldsymbol{u}} J \text{ subject to the constraint } \dot{\boldsymbol{x}} = \boldsymbol{F}(\boldsymbol{x}, \boldsymbol{u})$$

The strategy will be to assume that $\boldsymbol{u}$ maximizes J and then use this assumption to derive other conditions for a maximum. These arguments depend on making a perturbation in $\boldsymbol{u}$ and seeing what happens. Since $\boldsymbol{u}$ affects $\boldsymbol{x}$, the calculations become a little involved, but the argument is just a matter of careful bookkeeping. The main trick is to add additional terms to J that sum to zero. Let's start by appending the dynamic equation to J as before, but this time using continuous Lagrange multipliers $\boldsymbol{\lambda}(t)$:

$$\bar{J} = J - \int_0^T \boldsymbol{\lambda}^T [\dot{\boldsymbol{x}} - \boldsymbol{F}(\boldsymbol{x}, \boldsymbol{u})] dt$$

Hamiltonian

Anticipating what is about to happen, we define the *Hamiltonian* $H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})$ as

$$H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}) \equiv \boldsymbol{\lambda}^T[\boldsymbol{F}(\boldsymbol{x}, \boldsymbol{u})] + \ell(\boldsymbol{x}, \boldsymbol{u})$$

so that the expression for $\bar{J}$ becomes

$$\bar{J} = \psi[\boldsymbol{x}(T)] + \int_0^T [H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}) - \boldsymbol{\lambda}^T\dot{\boldsymbol{x}}]dt$$

δJ

Now let's examine the effects of a small change in $\boldsymbol{u}$, as shown in Figure 6 on $\bar{J}$, just keeping track of the change $\delta\bar{J}$:

$$\delta\bar{J} = \psi[\boldsymbol{x}(T)+\delta\boldsymbol{x}(T)] - \psi[\boldsymbol{x}(T)] + \int_0^T [H(\boldsymbol{\lambda}, \boldsymbol{x}+\delta\boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}) - \boldsymbol{\lambda}^T\delta\dot{\boldsymbol{x}}]dt$$

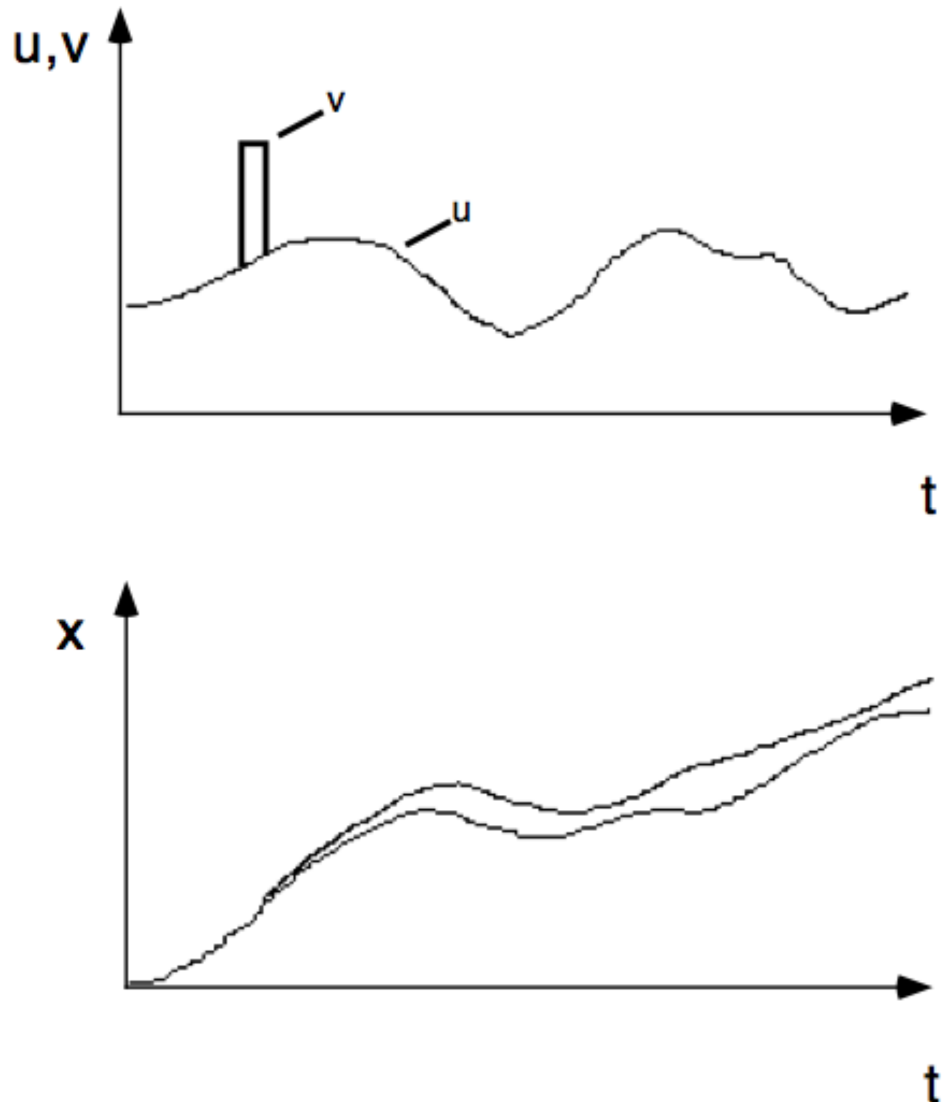

Integration by parts

Using the expression for integration by parts for $\int \boldsymbol{\lambda}^T \delta \dot{\boldsymbol{x}} dt$:

$$\int_0^T \boldsymbol{\lambda}^T \delta \dot{\boldsymbol{x}}\, dt = \boldsymbol{\lambda}^T(T)\delta \boldsymbol{x}(T) - \boldsymbol{\lambda}^T(0)\delta \boldsymbol{x}(0) - \int_0^T \dot{\boldsymbol{\lambda}}^T \delta \boldsymbol{x}\, dt$$

Now substitute this into the expression for $\delta \bar{J}$,

$$\begin{aligned}\delta \bar{J} \;&=\; \psi[\boldsymbol{x}(T) + \delta \boldsymbol{x}(T)] - \psi[\boldsymbol{x}(T)] - \boldsymbol{\lambda}(T)^T \delta \boldsymbol{x}(T) \\ &\quad + \int_0^T [H(\boldsymbol{\lambda}, \boldsymbol{x} + \delta \boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}) + \dot{\boldsymbol{\lambda}}^T \delta \boldsymbol{x}] dt\end{aligned}$$

Variational analysis

Now concentrate just on the first two terms in the integral:

$$\int_0^T [H(\boldsymbol{\lambda}, \boldsymbol{x} + \delta\boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})]dt$$

First add and subtract $H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v})$:

$$= \int_0^T [H(\boldsymbol{\lambda}, \boldsymbol{x} + \delta\boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) + H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})]dt$$

Next expand the first term inside the integral in a Taylor series and neglect terms above first order,

$$\cong \int_0^T (H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v})^T \delta\boldsymbol{x} + H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}))dt$$

where $H_{\boldsymbol{x}}$ is the partial $\frac{\partial H}{\partial \boldsymbol{x}}$, which is

$$\begin{pmatrix} H_{x_1} \\ \vdots \\ H_{x_n} \end{pmatrix}$$

Variational Analysis Cont.

Now add and subtract $H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})^T \delta \boldsymbol{x}$:

$$= \int_0^T \{H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})^T \delta \boldsymbol{x} + [H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})]^T \delta \boldsymbol{x} + H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})\} dt$$

The term $[H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})]^T \delta \boldsymbol{x}$ can be neglected because it is the product of two small terms, $[H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})]$ and $\delta \boldsymbol{x}$, and thus is a second-order term. Thus

$$\cong \int_0^T (H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}) \delta \boldsymbol{x} + H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})) dt$$

Adjoint Equation

Finally, substitute this expression back into the original equation for δJ, yielding

$$\begin{aligned}\delta \bar{J} \cong &\{\psi_{\boldsymbol{x}}[\boldsymbol{x}(T)] - \boldsymbol{\lambda}^T(T)\}\delta \boldsymbol{x}(T) \\ &+ \int_0^T [H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}) + \dot{\boldsymbol{\lambda}}^T]\delta \boldsymbol{x}\, dt \\ &+ \int_0^T [H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})]dt\end{aligned}$$

Since we have the freedom to pick $\boldsymbol{\lambda}$, just to make matters simpler, pick it so that the first integral vanishes:

$$-\dot{\boldsymbol{\lambda}}^T = H_{\boldsymbol{x}}(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})$$

$$\boldsymbol{\lambda}^T(T) = \psi_{\boldsymbol{x}}[\boldsymbol{x}(T)]$$

Condition for a maximum

Now all $\delta \bar{J}$ has left is

$$\delta \bar{J} = \int_0^T [H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) - H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u})] dt$$

From this equation it follows that the optimal control $\boldsymbol{u}^*$ must be such that

$$H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}^*) \geq H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}), \;\; \boldsymbol{u} \in U \tag{3}$$

To see this point, suppose that it were not true, that is, that for some interval of time there was a $\boldsymbol{v} \in U$ such that

$$H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{v}) \geq H(\boldsymbol{\lambda}, \boldsymbol{x}, \boldsymbol{u}^*)$$

This assumption would mean that you could adjust the integral so that the perturbation $\delta \bar{J}$ is positive, contradicting the original assumption that $\bar{J}$ is maximimized by $\boldsymbol{u}^*$. Therefore Equation 3 must hold.

Summary

In addition to the dynamic equations

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}, \boldsymbol{u})$$

and associated initial condition

$$\boldsymbol{x}(0) = \boldsymbol{x}_0$$

the Lagrange multipliers also must obey a constraint equation

$$-\dot{\boldsymbol{\lambda}}^T = H_{\boldsymbol{x}}$$

that has a *final condition*

$$\boldsymbol{\lambda}^T(T) = \psi_{\boldsymbol{x}}[\boldsymbol{x}(T)]$$

The equation for $\boldsymbol{\lambda}$ is known as the *adjoint equation*. In addition, for all t, the optimal control u is such that

$$H[\boldsymbol{\lambda}(t), \boldsymbol{x}(t), \boldsymbol{v}] \leq H[\boldsymbol{\lambda}(t), \boldsymbol{x}(t), \boldsymbol{u}(t)]$$

where H is the Hamiltonian

$$H = \boldsymbol{\lambda}^T f(\boldsymbol{x}, \boldsymbol{u}) + \ell(\boldsymbol{x}, \boldsymbol{u})$$

The Cart

The dynamic equation is

$$\ddot{x} = -\dot{x} + u(t)$$

with initial conditions

$$x(0) = 0$$
$$\dot{x}(0) = 0$$

The cost functional

$$J = x(T) - \frac{1}{2}\int_0^T u^2(t)dt$$

captures the desire to maximize the distance traveled in time T and at the same time penalize excessive accelerations.

Using the transformation of Section 5.2.1, the state variables x_1 and x_2 are defined by

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} x_2 \\ -x_2 + u \end{pmatrix}$$
$$x_1(0) = x_2(0) = 0$$
$$J = x_1(T) - \frac{1}{2}\int_0^T u^2 dt$$

The Hamiltonian is given by

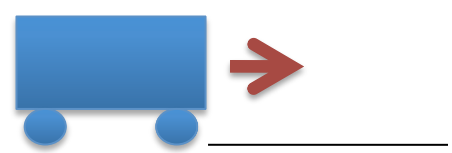

$$H = \lambda_1 x_2 - \lambda_2 x_2 + \lambda_2 u - \frac{1}{2}u^2$$

Solution

Differentiating this equation allows the determination of the adjoint system as

$$-\dot{\lambda}_1 = \frac{\partial H}{\partial x_1} = 0$$

$$-\dot{\lambda}_2 = \frac{\partial H}{\partial x_2} = \lambda_1 - \lambda_2$$

and its final condition can be determined from

$$\psi = x_1(T)$$

$$\boldsymbol{\lambda}(T) = \psi_{\boldsymbol{x}}(T) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

The simple form of the adjoint equations allows their direct solution. For λ_1,

$$\lambda_1 = \text{ const } = 1$$

For λ_2,we could use Laplace transform methods, but they have not been discussed, so lets make the incredible lucky guess:

$$\lambda_2 = 1 - e^{t-T}$$

For a maximum differentiate H with respect to u,

$$\frac{\partial H}{\partial u} = 0 \Rightarrow \lambda_2 - u = 0$$

$$u = \lambda_2 = 1 - e^{t-T}$$

Kalman Filter

Suppose you are able to make measurements z that are realated to a variable that you would like to know by a linear relationship:

$$z = Hx + \nu$$

The term v represents unwanted noise and is assumed to have the statistics

$$E(\nu) = 0$$

and

$$E(\nu\nu^T) = R$$

Somehow you have already estimated the variable x as $\bar{x}$.

Now you would like to use z to improve the estimate of x. A logical way to proceed would be to weight the two different sources of knowledge, z and $\bar{x}$.

$$J(x) = \frac{1}{2}[(x-\bar{x})^T M^{-1}(x-\bar{x})+(z-Hx)^T R^{-1}(z-Hx)]$$

Kalman Filter

Suppose you are able to make measurements z that are realated to a variable that you would like to know by a linear relationship:

$$z = Hx + \nu$$

The term v represents unwanted noise and is assumed to have the statistics

$$E(\nu) = 0$$

and

$$E(\nu\nu^T) = R$$

Somehow you have already estimated the variable x as $\bar{x}$.

Now you would like to use z to improve the estimate of x. A logical way to proceed would be to weight the two different sources of knowledge, z and $\bar{x}$.

$$J(x) = \frac{1}{2}[(x-\bar{x})^T M^{-1}(x-\bar{x})+(z-Hx)^T R^{-1}(z-Hx)]$$

At the minimum,

$$dJ = 0 = dx^T[M^{-1}(x - \bar{x}) - H^T R^{-1}(z - Hx)]$$

To satisfy this in general, the term in [] must be zero

$$M^{-1}(\hat{x} - \bar{x}) - H^T R^{-1}(z - H\hat{x}) = 0$$

$$(M^{-1} - H^T R^{-1} H)\hat{x} = M^{-1}\bar{x} + H^T R^{-1} z$$

Adding and subtracting $H^T R^{-1} H\bar{x}$ to the RHS,

$$(M^{-1} - H^T R^{-1} H)\hat{x} = (M^{-1} - H^T R^{-1} H)\bar{x} + H^T R^{-1}(z - H\hat{x})$$

Kalman Filter

Now use

$$P = M^{-1} - H^T R^{-1} H$$

to write the estimate for $\hat{x}$ as

$$\hat{x} = \bar{x} + P^{-1} H^T R^{-1} (z - H\bar{x})$$

This is the *least squares estimate* for x. It can be shown that P is the covariance of the error in the estimate, that is

$$P = E[(\hat{x} - x)(\hat{x} - x)^T]$$

Kalman Filter

Now suppose that you would like to estimate the value of a variable as before. However now this variable x_1 is related to a previous variable x_0 by

$$x_1 = Ax_0 + B\mu$$

where

$$E[x_0] = \bar{x}_0$$

$$E[\mu] = \bar{\mu}$$

$$E[(\mu - \bar{\mu}_0)(\mu - \bar{\mu}_0)^T] = Q_0$$

and

$$E[(\hat{x}_0 - x_0)(\hat{x}_0 - x_0)^T] = P_0$$

Kalman Filter

This problem is easy since we have just solved it! The solution is given by

$$\hat{x}_1 = \bar{x}_1 + P_1^{-1} H^T R^{-1} (z_1 - H\bar{x}_1)$$

where

$$P_1 = M_1^{-1} - H^T R^{-1} H$$

The new wrinkle is that now

$$\bar{x}_1 = A\bar{x}_0 + B\bar{\mu}$$

and

$$M_1 = AP_0 A^T + BQ_0 B^T$$

This can be iteratively extended to handle the case where

$$x_{k+1} = Ax_k + B\mu$$