Optimal feedback control

CoSMo 2017)
Gunnar Blohm

Outline

- Inverted pendulum
  - Physics and simulations (hands-on)
  - Compute cost
- Control
  - Brute force optimization - fitting
  - Optimal control
- Kalman filtering
  - Process and measurement noise
  - Optimal estimation

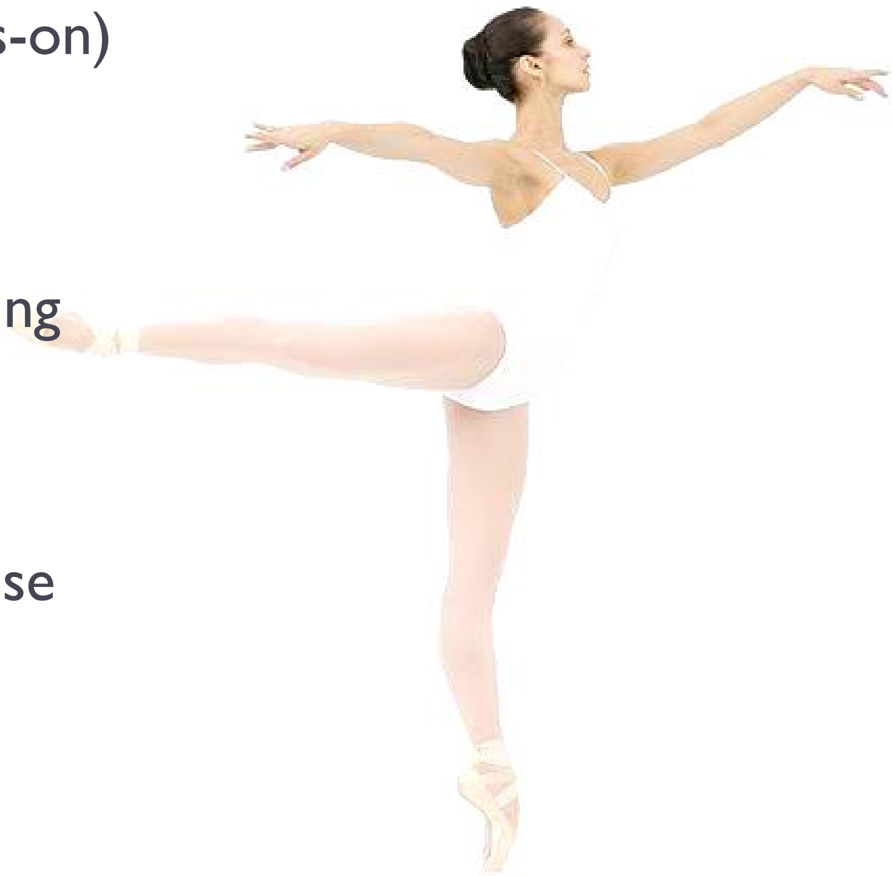

Inverted pendulum

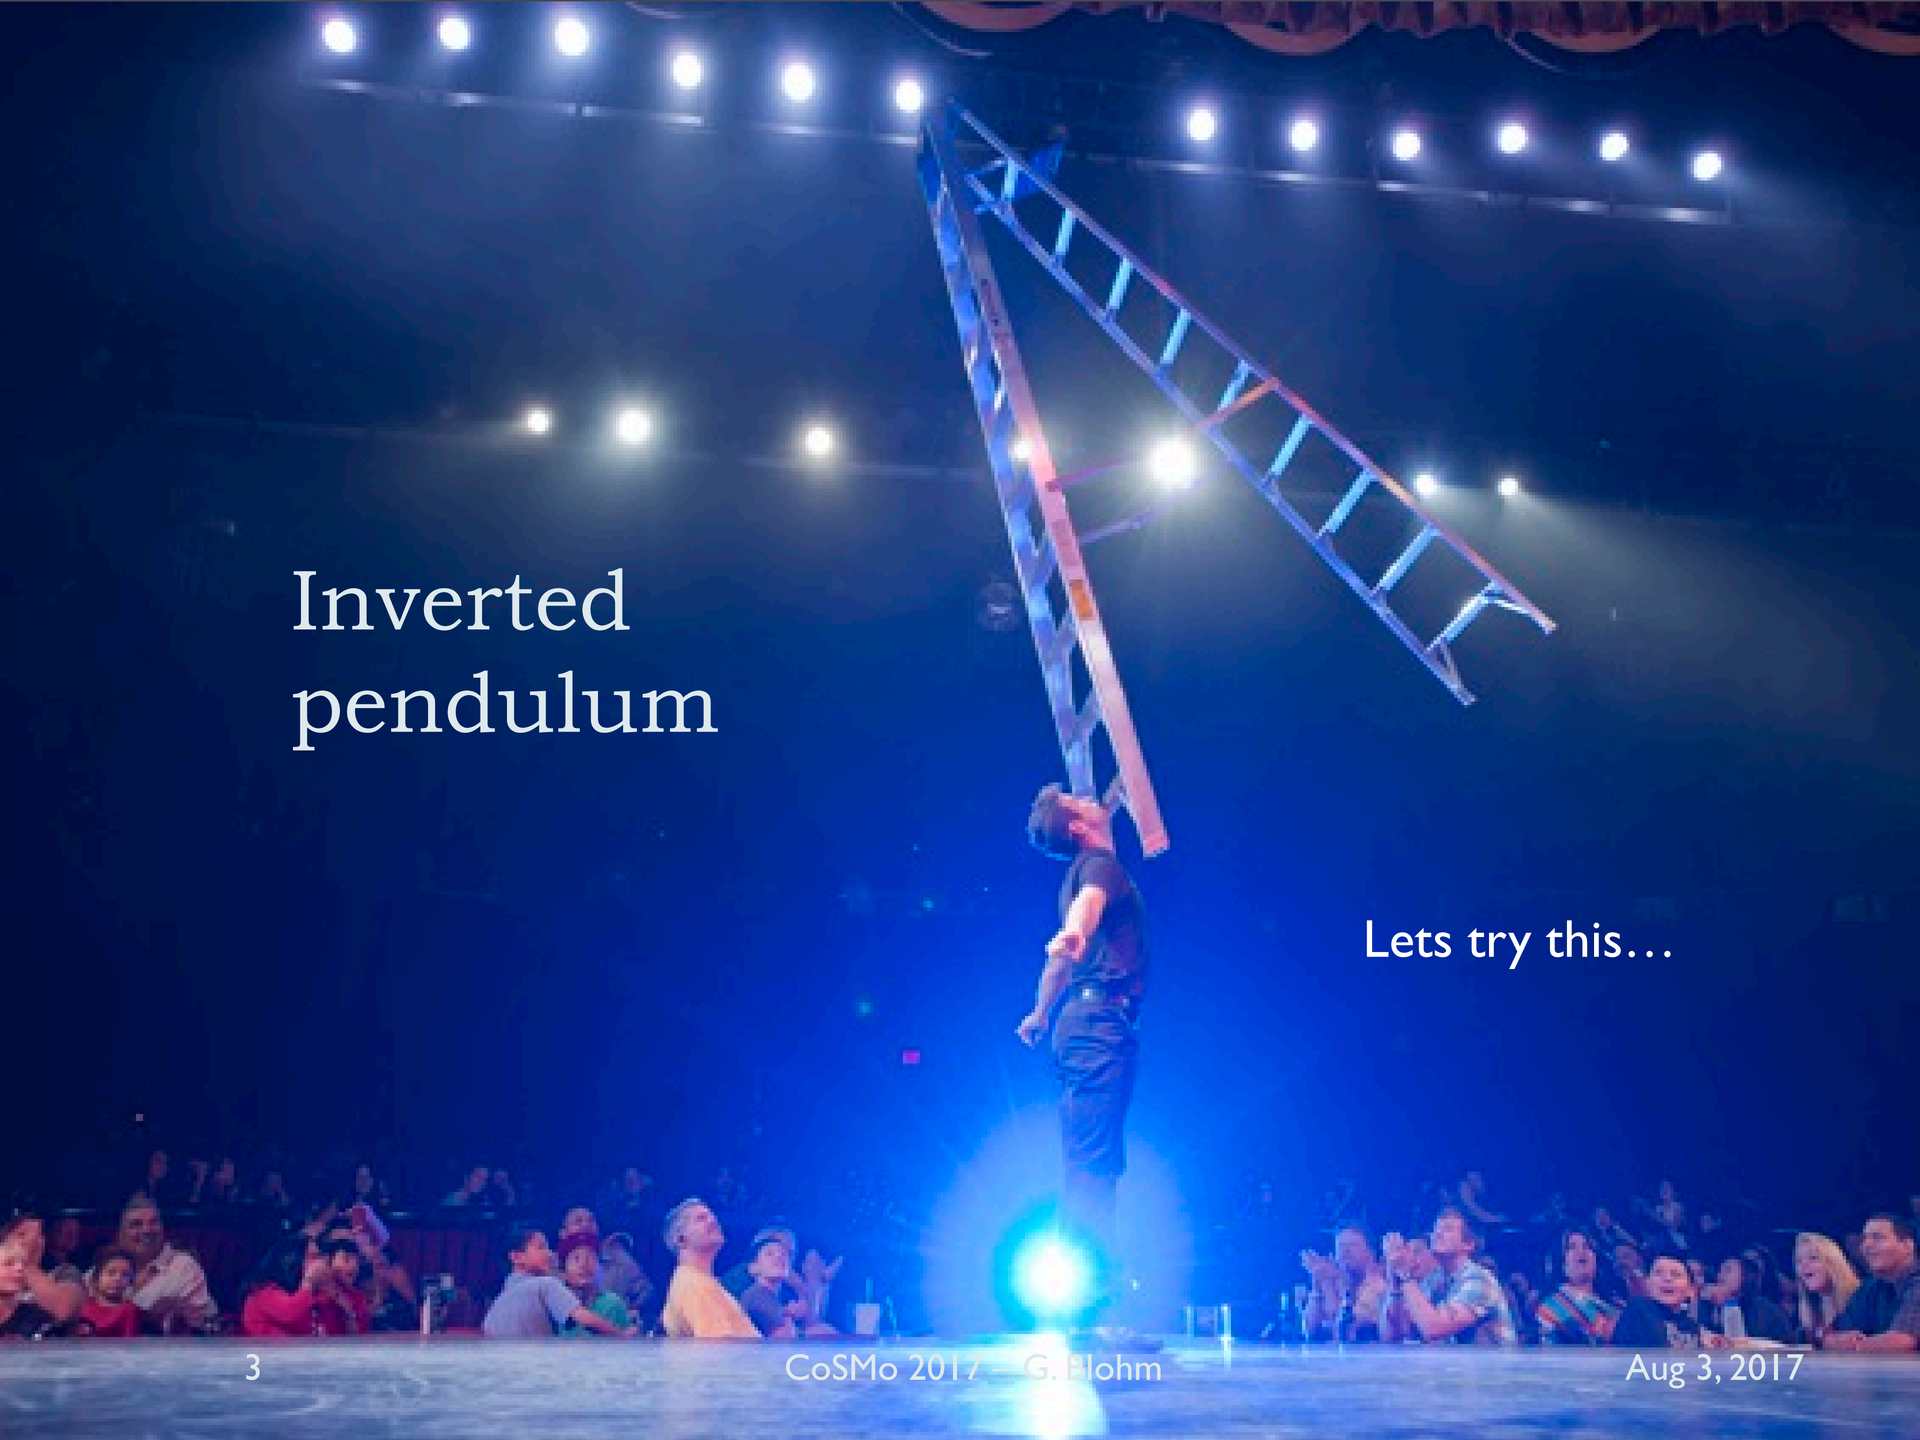

Lets try this…

The inverted pendulum

- Motorized cart
- Goal: bring cart from A to B while maintaining verticality
- Inherently unstable

- Real-world examples:
  - Attitude control of booster rocket at takeoff
  - Balancing a latter on your chin…

- Specifics
  - (M) mass of the cart = 0.5 kg
  - (m) mass of the pendulum = 0.2 kg
  - (b) coefficient of friction for cart = 0.1 N/m/sec
  - (l) length to pendulum center of mass = 0.3 m
  - (I) mass moment of inertia of the pendulum = 0.006 kg.m^2
  - (F) force applied to the cart(x) cart position coordinate
  - (theta) pendulum angle from vertical (down)

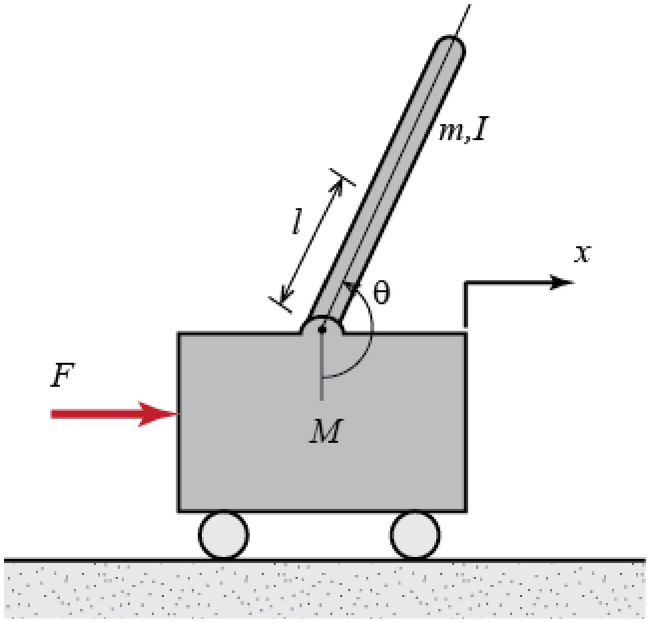

The inverted pendulum

- Equations of motion
  - Forces on cart:

    $$M\ddot{x} + b\dot{x} + N = F$$

  - Horizontal forces on pendulum:

    $$N = m\ddot{x} + ml\ddot{\theta}\cos\theta - ml\dot{\theta}^2\sin\theta$$

  - Forces perpendicular to pendulum:

    $$P\sin\theta + N\cos\theta - mg\sin\theta = ml\ddot{\theta} + m\ddot{x}\cos\theta$$

  - Moments at centroid of pendulum:

    $$-Pl\sin\theta - Nl\cos\theta = I\ddot{\theta}$$

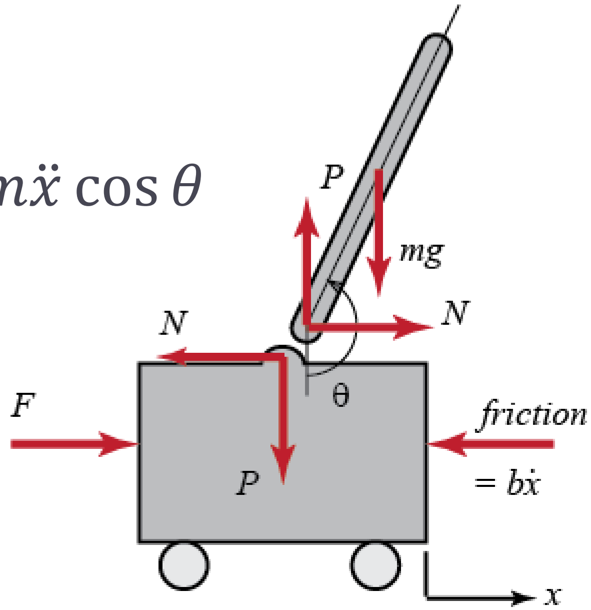

The inverted pendulum

- Simplifications
  - Point-mass and mass-less pole: I = 0
  - No friction: b = 0

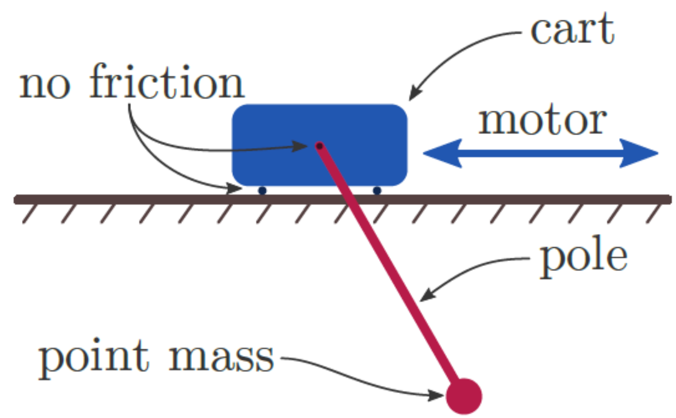

- $$\begin{pmatrix} \cos\theta & l \\ m+M & ml\cos\theta \end{pmatrix} \begin{pmatrix} \ddot{x} \\ \ddot{\theta} \end{pmatrix} = \begin{pmatrix} -g\sin\theta \\ F + ml\dot{\theta}^2\sin\theta \end{pmatrix}$$
  - Simulate system behavior without control
  - Experiment with initial conditions…

Optimal control

The problem

- The body can be considered as a dynamical system

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

- Applying control *u* in a fully deterministic system will result in a trajectory

$$x(t) = x_0 + \int_{t_0}^{t} f(x(s), u(s))\, ds.$$

- But how can we find *u* to steer the system towards a specific goal?

Crevecoeur, Cluff, Scott (2014)

Control solution

- What is the **best** way to steer the system?

- Insight (Pontryagin & Bellman, 1950s): a good controller should achieve a certain optimality principle!
  - Cost function

- Reformulation of control system
  - Minimize cost
  - Constraint: system dynamics

The "stupid" solution

- Brute force optimization!

- Select an arbitrary function $u = f(x, p)$
  - Parameters p
  - Find the optimal parameter set p' that minimizes your error

Optimal control cont’d

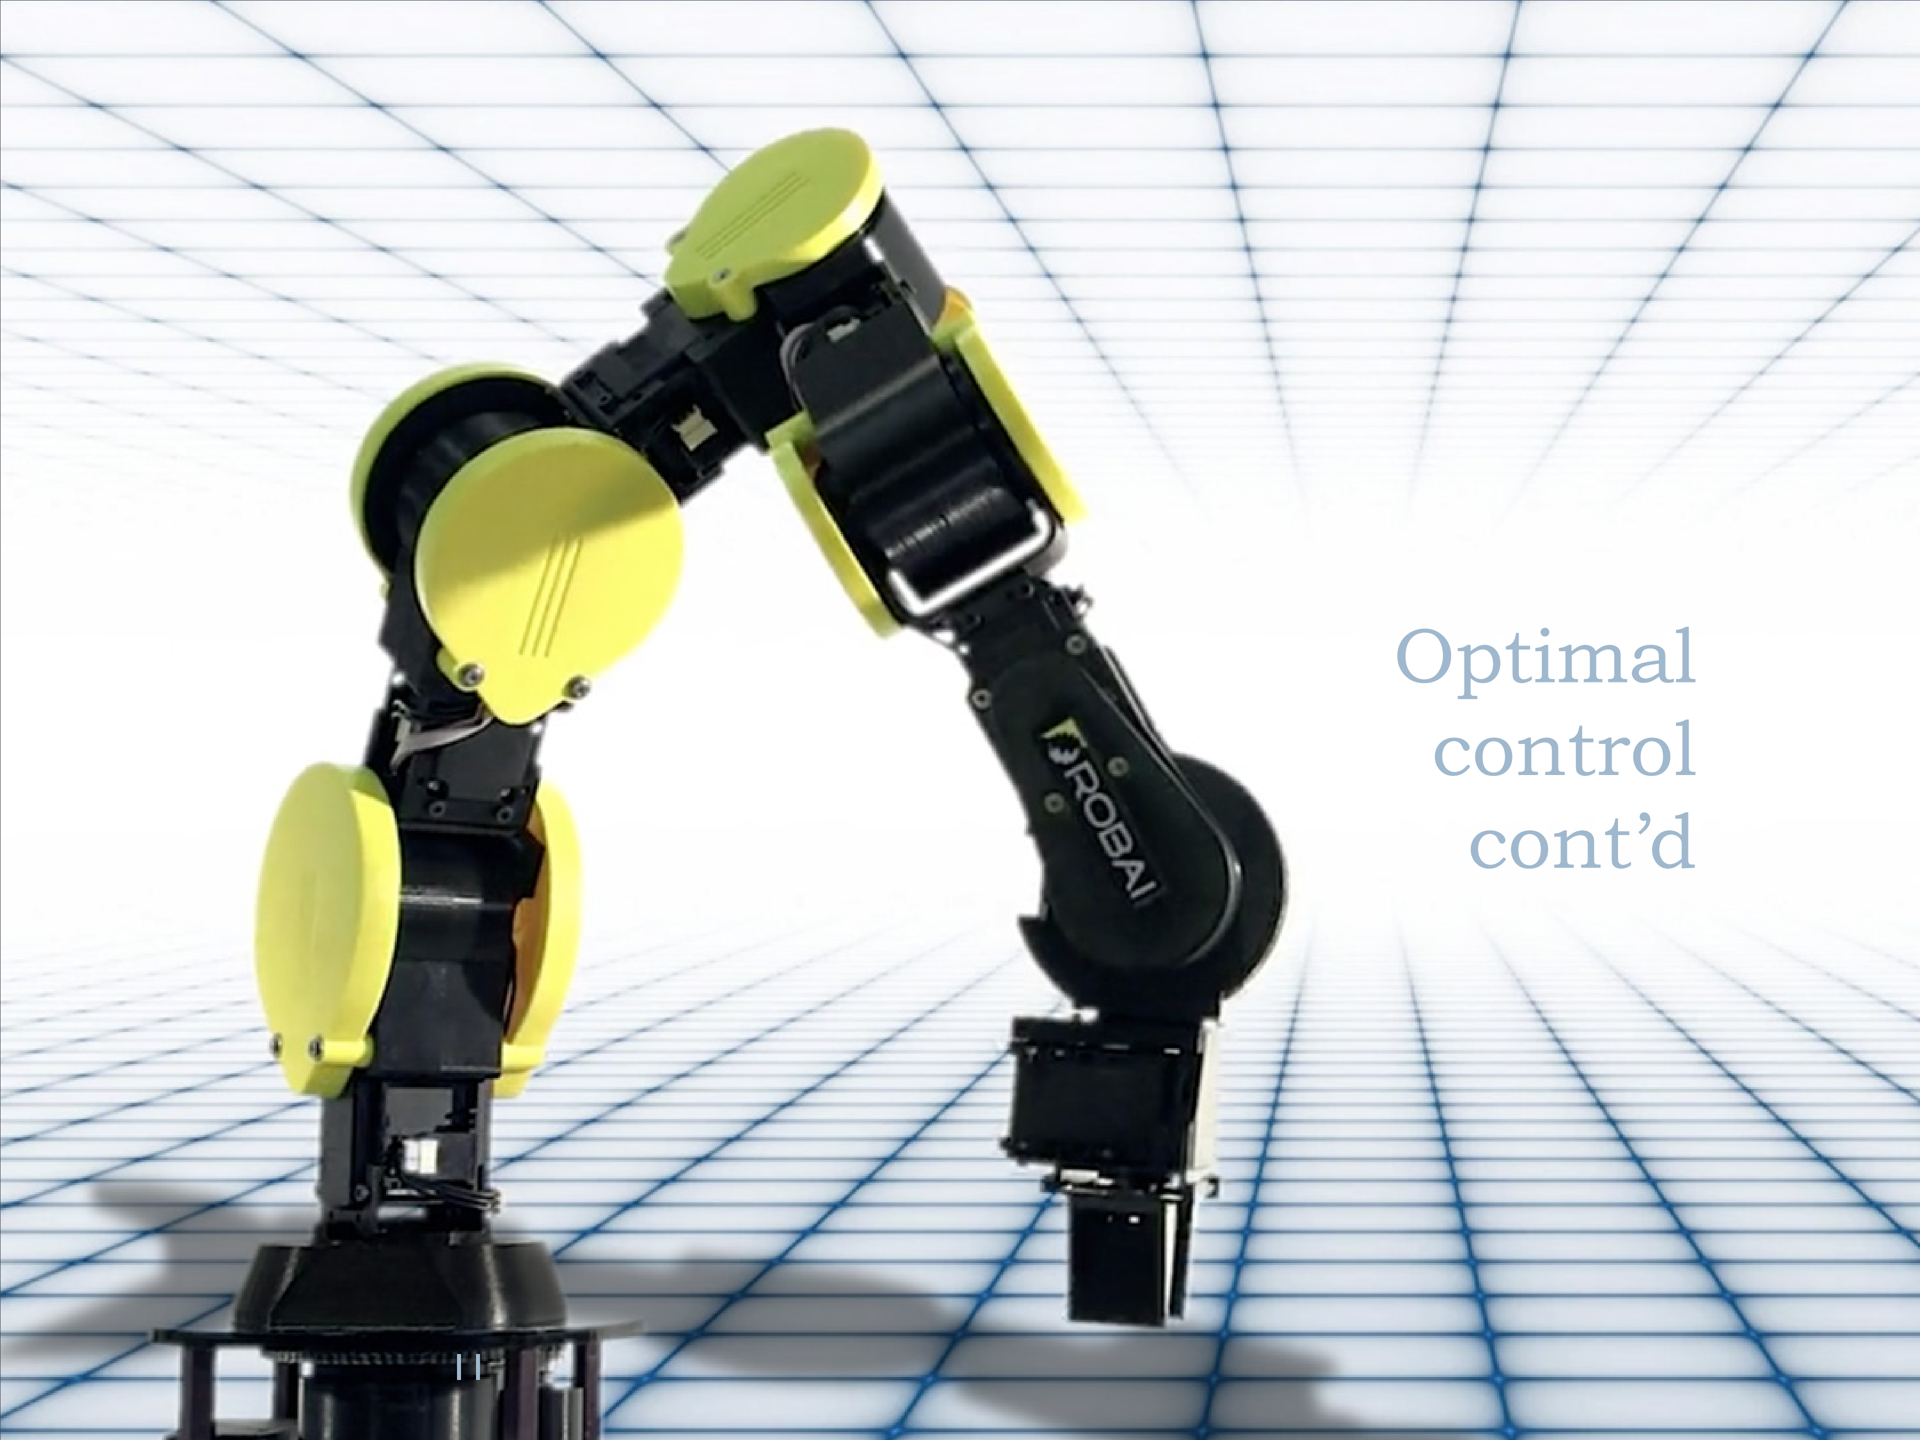

The inverted pendulum

- Approximations
  - We will apply LINEAR control methods
    - Thus we need to linearize any non-linear terms
  - We assume control close to equilibrium π
  - $\cos\theta = \cos(\pi - \phi) \approx -1$
  - $\sin\theta = \sin(\pi - \phi) \approx -\phi$
  - $\dot{\theta}^2 = \dot{\phi}^2 \approx 0$

  - Solution:
  - $(I + ml^2)\ddot{\phi} - mgl\phi = ml\ddot{x}$
  - $(M + m)\ddot{x} + b\dot{x} - ml\ddot{\phi} = F = u$ ← control

The inverted pendulum

- State-space formulation

$$\begin{bmatrix} \dot{x} \\ \ddot{x} \\ \dot{\phi} \\ \ddot{\phi} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & \frac{-(I+ml^2)b}{I(M+m)+Mml^2} & \frac{m^2gl^2}{I(M+m)+Mml^2} & 0 \\ 0 & 0 & 0 & 1 \\ 0 & \frac{-mlb}{I(M+m)+Mml^2} & \frac{mgl(M+m)}{I(M+m)+Mml^2} & 0 \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \\ \phi \\ \dot{\phi} \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{I+ml^2}{I(M+m)+Mml^2} \\ 0 \\ \frac{ml}{I(M+m)+Mml^2} \end{bmatrix} u$$

A **B**

$$\mathbf{y} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \\ \phi \\ \dot{\phi} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix} u$$

C **D**

Noiseless optimal solution

- Defining an objective function, i.e. cost:

$$J(x,u) = g(x(t_f)) + \int_{t_0}^{t_f} L(x(t),u(t))dt.$$

- Optimal control = find *u* that minimizes *J(x,u)*
- In that case, the Hamilton-Jacobi-Bellman equation is satisfied for the cost-to-go *J(x,u)*

$$-\frac{\partial J(x,u)}{\partial t} = \min_{u \in U}\left\{L(x,u) + \frac{\partial J(x,u)}{\partial x} f(x,u)\right\}.$$

  - Trade-off between instantaneous cost and gradient of cost-to-go

Crevecoeur, Cluff, Scott (2014)

What about noise?

- Dynamical system equation with noise

$$dX = F(X,u)dt + G(X,u)dW$$

System dynamics — Random disturbances (Brownian motion)

- Only the expected outcome can be optimized, but the real outcome will deviate because of noise

$$J_t(X_t,u_t) = \min_u \left\{ L(X_t,u_t) + E\left[J_{t+1}(X_{t+1},u_{t+1}) \mid X_t,u_t \right] \right\}$$

Current (running) cost — Expected future cost-to-go

Crevecoeur, Cluff, Scott (2014)

Linear Quadratic Gaussian regulator (LQG)

  - Linear dynamics
  - Quadratic cost function
  - Gaussian noise
- Simplest case with analytical solution
  - Minimize J under the constraint of the system dynamics
  - Requires Lagrange multipliers…
  - E.g. maximize $f(x, y)$ subject to $g(x, y) = c$
    - Lagrangian: $L(x,y,\lambda) = f(x, y) + \lambda\,(g(x, y) - c)$
    - Solve: $\nabla_{x,y,\lambda} L = 0$
  - For LQG: one λ for each time step → iterative procedure

LQ(G) solution: control

- Let's consider a fully observable state

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t$$

$$J(\mathbf{x},\mathbf{u}) = \mathbf{x}_n^{\mathrm{T}}\mathbf{Q}_n\mathbf{x}_n + \sum_{t=0}^{n-1}\left(\mathbf{x}_t^{\mathrm{T}}\mathbf{Q}_t\mathbf{x}_t + \mathbf{u}_t^{\mathrm{T}}\mathbf{R}_t\mathbf{u}_t\right)$$

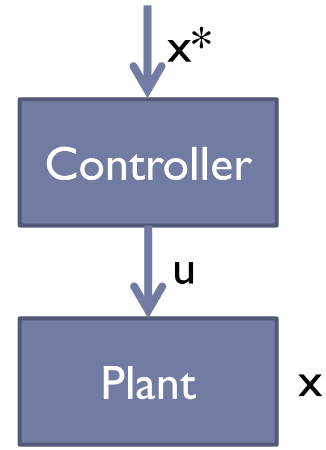

- Lagrange function

$$L(\mathbf{x},\mathbf{u},\lambda) = \mathbf{x}_n^{\mathrm{T}}\mathbf{Q}_n\mathbf{x}_n + \sum^{n-1}\mathbf{x}_t^{\mathrm{T}}\mathbf{Q}_t\mathbf{x}_t + \mathbf{u}_t^{\mathrm{T}}\mathbf{R}_t\mathbf{u}_t - (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t - \mathbf{x}_{t+1})^{\mathrm{T}}\lambda_{t+1}$$

$$L(\mathbf{x},\mathbf{u},\lambda) = \mathbf{x}_n^{\mathrm{T}}\mathbf{Q}_n\mathbf{x}_n - \mathbf{x}_n^{\mathrm{T}}\lambda_n + \mathbf{x}_0^{\mathrm{T}}\lambda_0 + \sum_{t=0}^{n-1}\mathbf{x}_t^{\mathrm{T}}\mathbf{Q}_t\mathbf{x}_t + \mathbf{u}_t^{\mathrm{T}}\mathbf{R}_t\mathbf{u}_t - (\mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t)^{\mathrm{T}}\lambda_{t+1} - \mathbf{x}_t^{\mathrm{T}}\lambda_t$$

LQ(G) solution: control

- Now we take the gradients of L

$$\frac{\partial L}{\partial \mathbf{x}_n} = \mathbf{Q}_n \mathbf{x}_n - \lambda_n = 0 \longrightarrow \lambda_n = \mathbf{Q}_n \mathbf{x}_n \longrightarrow \lambda_t = \mathbf{S}_t \mathbf{x}_t$$

$$\frac{\partial L}{\partial \mathbf{u}_t} = \mathbf{R}_t \mathbf{u}_t - \mathbf{B}^{\mathrm{T}} \lambda_{t+1} = 0 \longrightarrow \mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}_{t+1}\mathbf{x}_{t+1}$$

(modified state equation)

$$\frac{\partial L}{\partial \lambda_{t+1}} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t - \mathbf{x}_{t+1} = 0$$

$$\mathbf{x}_{t+1} = (\mathbf{I} - \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}_{t+1})^{-1}\mathbf{A}\mathbf{x}_t$$

$$\frac{\partial L}{\partial \mathbf{x}_t} = \mathbf{Q}_t \mathbf{x}_t - \mathbf{A}^{\mathrm{T}} \lambda_{t+1} - \lambda_t = 0 \longrightarrow \mathbf{S}_t \mathbf{x}_t = \mathbf{Q}_t \mathbf{x}_t - \mathbf{A}^{\mathrm{T}}\mathbf{S}_{t+1}(\mathbf{I} - \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}_{t+1})^{-1}\mathbf{A}\mathbf{x}_t$$

$$\mathbf{S}_t = \mathbf{Q}_t - \mathbf{A}^{\mathrm{T}}\mathbf{S}_{t+1}(\mathbf{I} - \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}_{t+1})^{-1}\mathbf{A}$$

LQ(G) solution: control

- Since $\mathbf{S}_n = \mathbf{Q}_n$ we can run backwards in time to compute $\mathbf{S}_{n-1}$ and thus all $\mathbf{S}_t$ (cost to go gain)

$$\mathbf{S}_t = \mathbf{Q}_t - \mathbf{A}^{\mathrm{T}}\mathbf{S}_{t+1}(\mathbf{I} - \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}_{t+1})^{-1}\mathbf{A}$$

- Then we can run forward in time to find $\mathbf{x}_t$ and thus also λ_t and $\mathbf{u}_t$

$$\mathbf{x}_{t+1} = (\mathbf{I} - \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}_{t+1})^{-1}\mathbf{A}\mathbf{x}_t$$

$$\lambda_t = \mathbf{S}_t\mathbf{x}_t$$

$$\mathbf{u}_t = \mathbf{R}^{-1}\mathbf{B}^{\mathrm{T}}\mathbf{S}_{t+1}\mathbf{x}_{t+1}$$

- This is the solution without feedback…

Implementation of control

- Step 1: write down movement equations
- Step 2: transform into matrix format
  - Define "state"
  - Discretize
  - Write as $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t$
- Step 3: Define costs Q and R
- Step 4: Implement backward recursion for S (first loop)
- Step 5: Implement forward recursion for u…

LQ(G) solution: estimation

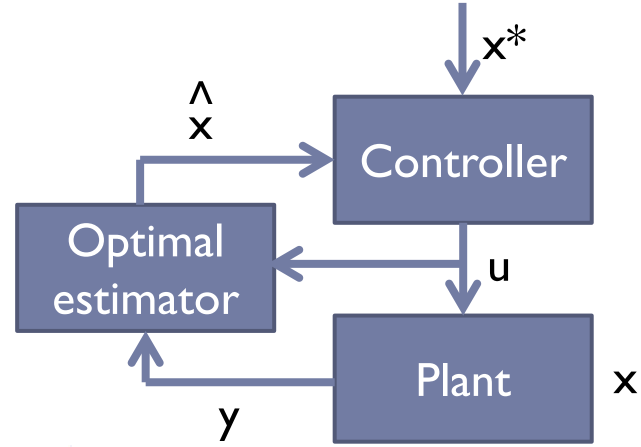

- Control system with noise

$$x_{k+1} = Ax_k + Bu_k + \xi_k$$
$$y_k = Hx_k + \omega_k.$$

- Bayes: combine prior and feedback

$$\hat{x}_{k+1} = (1-K) \times \text{prior} + K \times \text{feedback}$$
$$\hat{x}_{k+1} = A\hat{x}_k + Bu_k + K(y_k - H\hat{x}_k)$$

- Estimation error dynamics

$$e_{k+1} = (A - K_k H)e_k + \xi_k - K_k\omega_k$$

- Goal: optimize (Kalman) gain K to minimize error

$$K_k = \arg\min_K \| e_{k+1} \|^2$$

LQG solution: estimation & implementation

- Optimal estimates and Kalman gains are obtained in forward recursion

$$
\begin{aligned}
\hat{x}_{k+1} &= A\hat{x}_k + Bu_k + K(y_k - H\hat{x}_k), \\
K_k &= A\Sigma_k H^T (H\Sigma_k H^T + \Omega_\omega)^{-1}, \\
\Sigma_{k+1} &= \Omega_\xi + (A - K_k H)\Sigma_k A^T.
\end{aligned}
$$

  - Σ_1 known

- Step 6: include optimal estimation to complete LQG

Further considerations

- Control and estimation were solved separately but are valid together!
  - Separation principle!
  - Full control system:

$$\begin{bmatrix} x_{k+1} \\ e_{k+1} \end{bmatrix} = \begin{bmatrix} A - BL_k & BL_k \\ 0 & A - K_k H \end{bmatrix} \begin{bmatrix} x_k \\ e_k \end{bmatrix} + \begin{bmatrix} \xi_k \\ \xi_k - K_k \omega_k \end{bmatrix}$$

- Non-linear systems
  - Might be possible to linearize around current state using Taylor expansion…


That's all Folks!

PID control

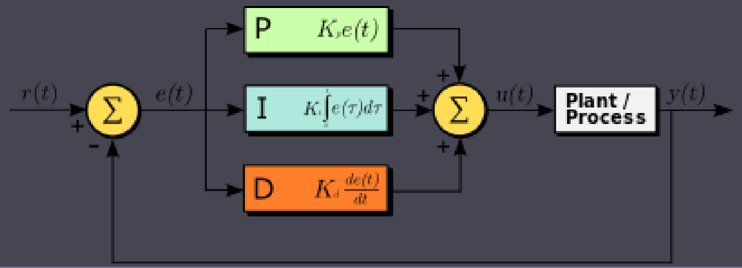

P
$K_p e(t)$
I
$K_i \int_0^t e(\tau) d\tau$
D
$K_d \frac{de(t)}{dt}$
$r(t)$
$e(t)$
$u(t)$
$y(t)$
Plant / Process

PID idea

- Observe plant
- Derive best action based on observation
- Use only simple operations
  - Proportional
  - Integral
  - Derivative
- Hope for the best…

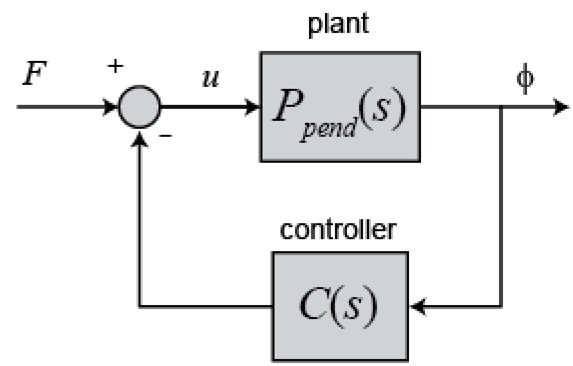

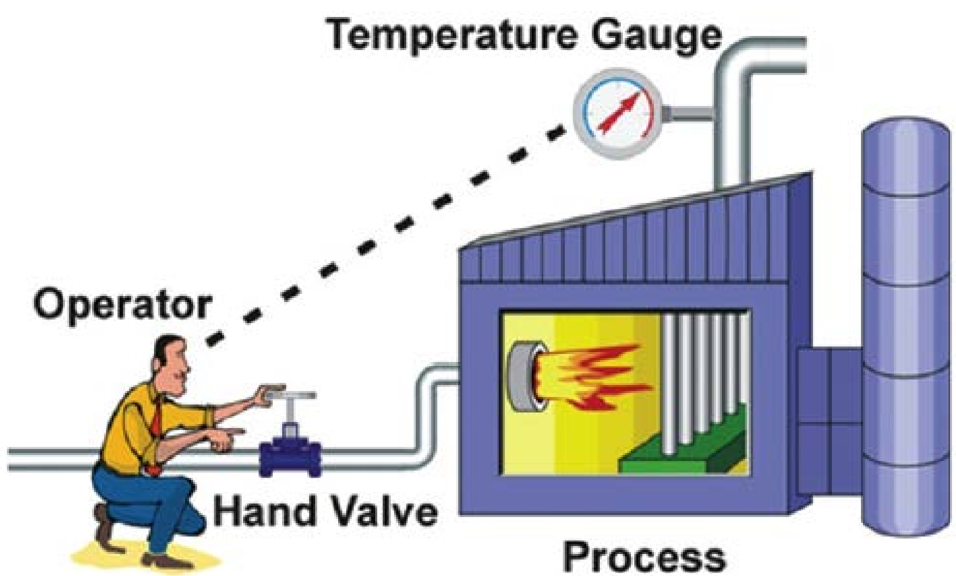

PID idea

- Proportional: increase speed
- Integral: eliminate steady-state errors
- Differential: decrease over-shoot & settling time

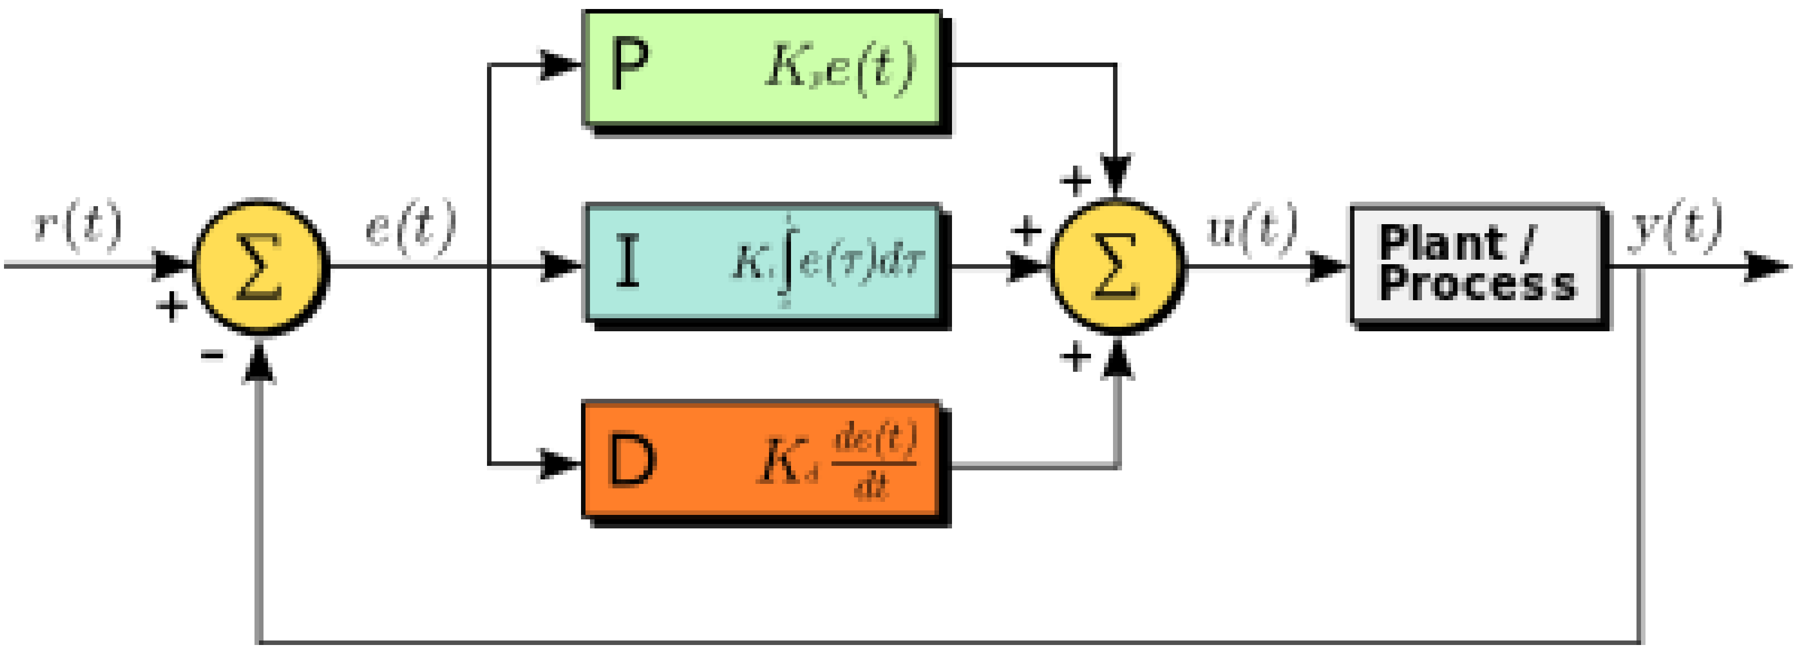

Uncertainty & optimal feedback control

Optimal state estimation

Neural control of movement

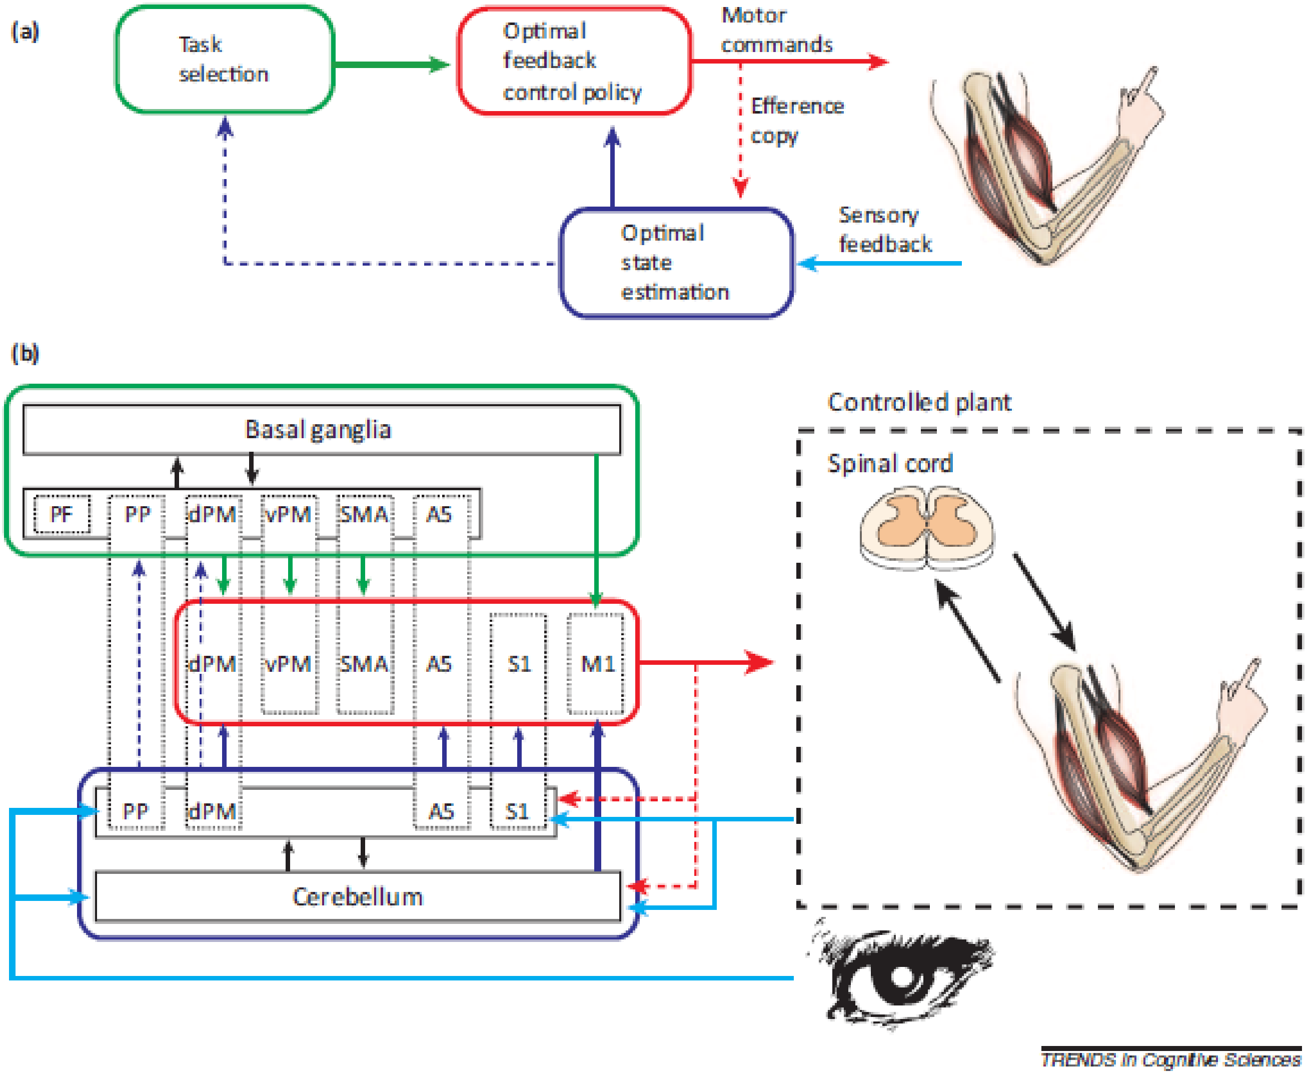

Scott, 2012

Optimal Feedback Control

- Todorov and Jordan, 2002

<u>Task Selection (Reaching)</u>
Target Position
Initial Arm Position
Nominal Speed

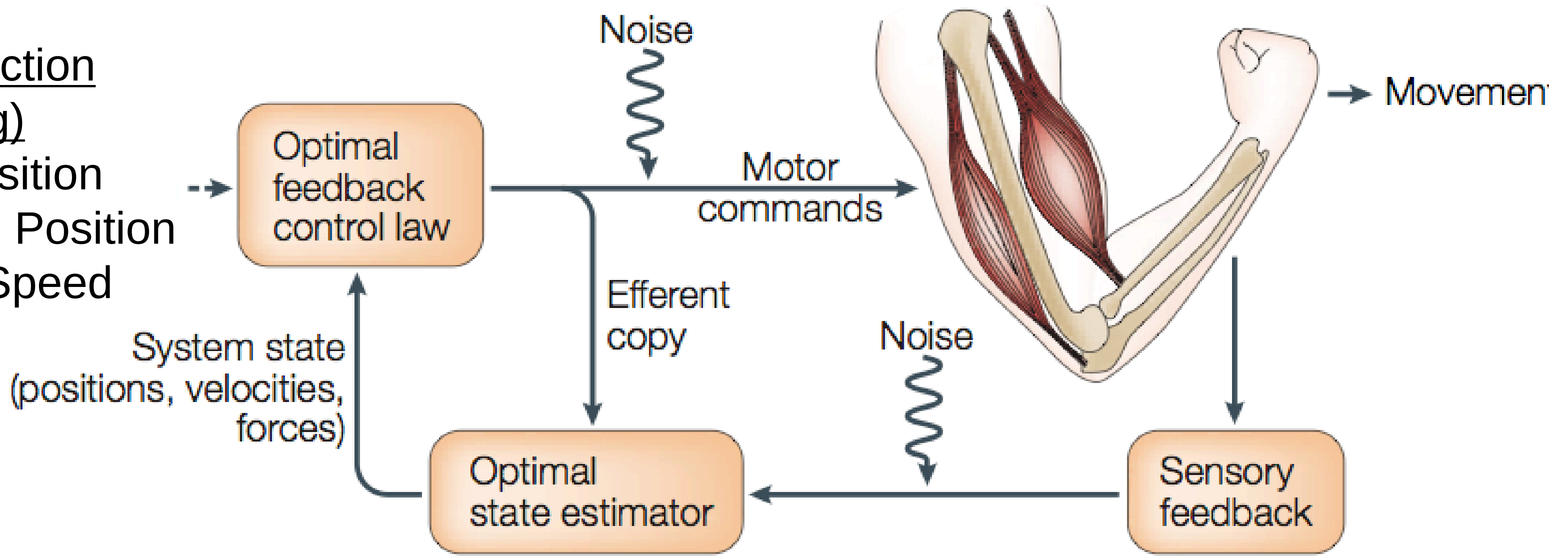

Scott, 2004

LQ(G) solution: estimation

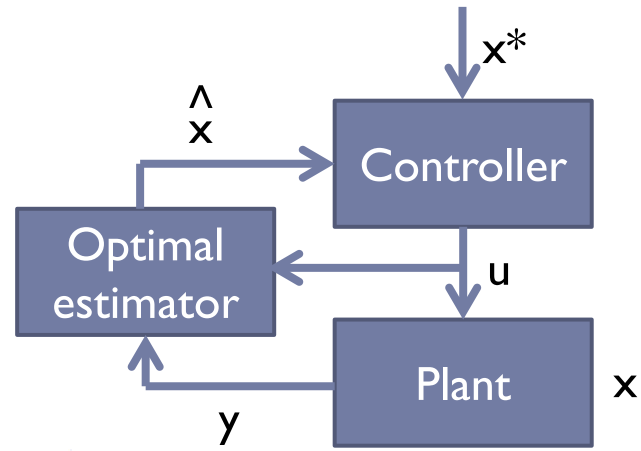

- Control system with noise

$$x_{k+1} = Ax_k + Bu_k + \xi_k$$
$$y_k = Hx_k + \omega_k.$$

- Bayes: combine prior and feedback

$$\hat{x}_{k+1} = (1-K) \times \text{prior} + K \times \text{feedback}$$
$$\hat{x}_{k+1} = A\hat{x}_k + Bu_k + K(y_k - H\hat{x}_k)$$

- Estimation error dynamics

$$e_{k+1} = (A - K_k H)e_k + \xi_k - K_k \omega_k$$

- Goal: optimize (Kalman) gain K to minimize error

$$K_k = \arg\min_K \| e_{k+1} \|^2$$

LQ(G) solution: estimation

- Minimization of error

$$\begin{aligned} K_k &= \arg\min_K \| e_{k+1} \|^2, \\ &= \arg\min_K \left[\mathrm{tr}\left(E(e_{k+1} e_{k+1}^T) \right) \right] \end{aligned}$$

- Only consider terms of the error covariance matrix that contain K

$$a(K_k) := \mathrm{tr}\left(-2K_k H \Sigma_k + K_k (H \Sigma_k H^T + \Omega_\omega) K_k^T \right)$$

- Minimization: $\nabla a(K_k) = 0$

$$\Rightarrow K_k = A \Sigma_k H^T (H \Sigma_k H^T + \Omega_\omega)^{-1}$$

LQ(G) solution: estimation

- Cheat sheet

$$\begin{aligned} x_{k+1} &= Ax_k + \xi_k \qquad \xi_k \sim N(0, \Omega_\xi) \\ y_k &= Hx_k + \omega_k \qquad \omega_k \sim N(0, \Omega_\omega) \end{aligned}$$

$$x_{k+1} \sim N(\mu_{t+1}, \Sigma_{t+1})$$

- From Kalman filter theory

$$\begin{aligned} \Sigma_{k+1|k} &= A\Sigma_k A^T + \Omega_\xi, \\ K_{k+1} &= \Sigma_{k+1|k} H^T \left(H\Sigma_{k+1|k} H^T + \Omega_\omega \right)^{-1} \\ \mu_{k+1} &= A\mu_k + K_{k+1}(y_{k+1} - HA\mu_k) \\ \Sigma_{k+1} &= (I - K_{k+1}H)\Sigma_{k+1|k}. \end{aligned}$$

LQG solution: estimation & implementation

- Optimal estimates and Kalman gains are obtained in forward recursion

$$\begin{aligned}\hat{x}_{k+1} &= A\hat{x}_k + Bu_k + K(y_k - H\hat{x}_k),\\ K_k &= A\Sigma_k H^T (H\Sigma_k H^T + \Omega_\omega)^{-1},\\ \Sigma_{k+1} &= \Omega_\xi + (A - K_k H)\Sigma_k A^T.\end{aligned}$$

  - Σ_1 known

- Step 6: include optimal estimation to complete LQG

Further considerations

- Control and estimation were solved separately but are valid together!
  - Separation principle!
  - Full control system:

$$\begin{bmatrix} x_{k+1} \\ e_{k+1} \end{bmatrix} = \begin{bmatrix} A - BL_k & BL_k \\ 0 & A - K_k H \end{bmatrix} \begin{bmatrix} x_k \\ e_k \end{bmatrix} + \begin{bmatrix} \xi_k \\ \xi_k - K_k \omega_k \end{bmatrix}$$

- Non-linear systems
  - Might be possible to linearize around current state using Taylor expansion…

Optimal feedback control

- Key feature: errors are only corrected if they affect the goal, otherwise they are ignored
  - Minimum intervention principle

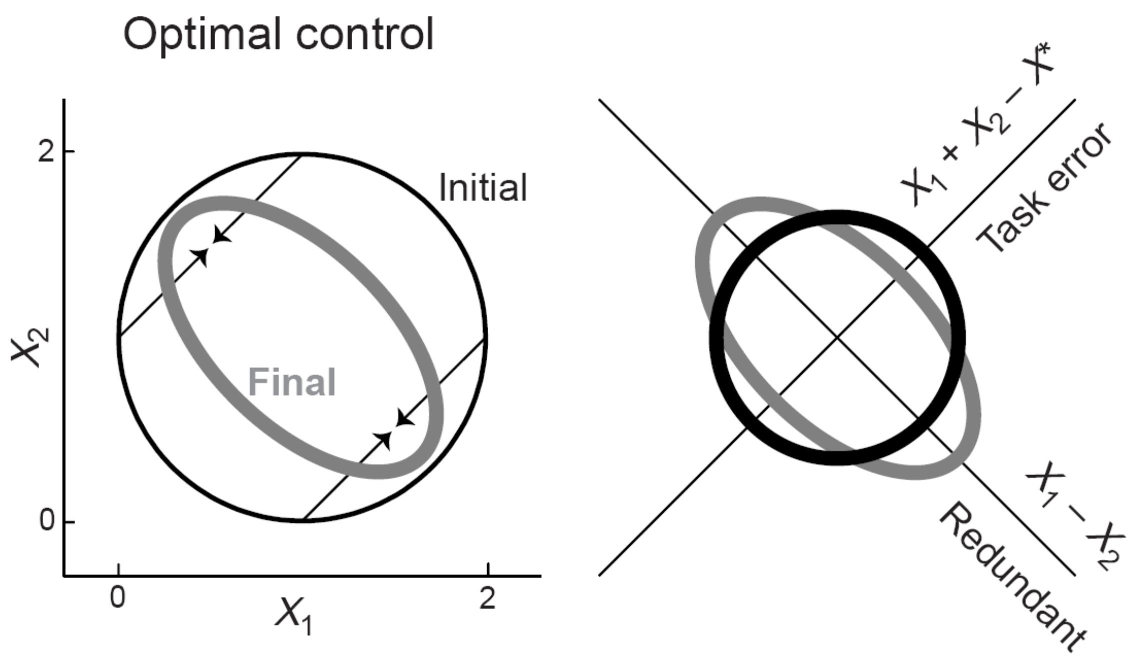

Torodov & Jordan, 2002

Minimum intervention principle: behaviour

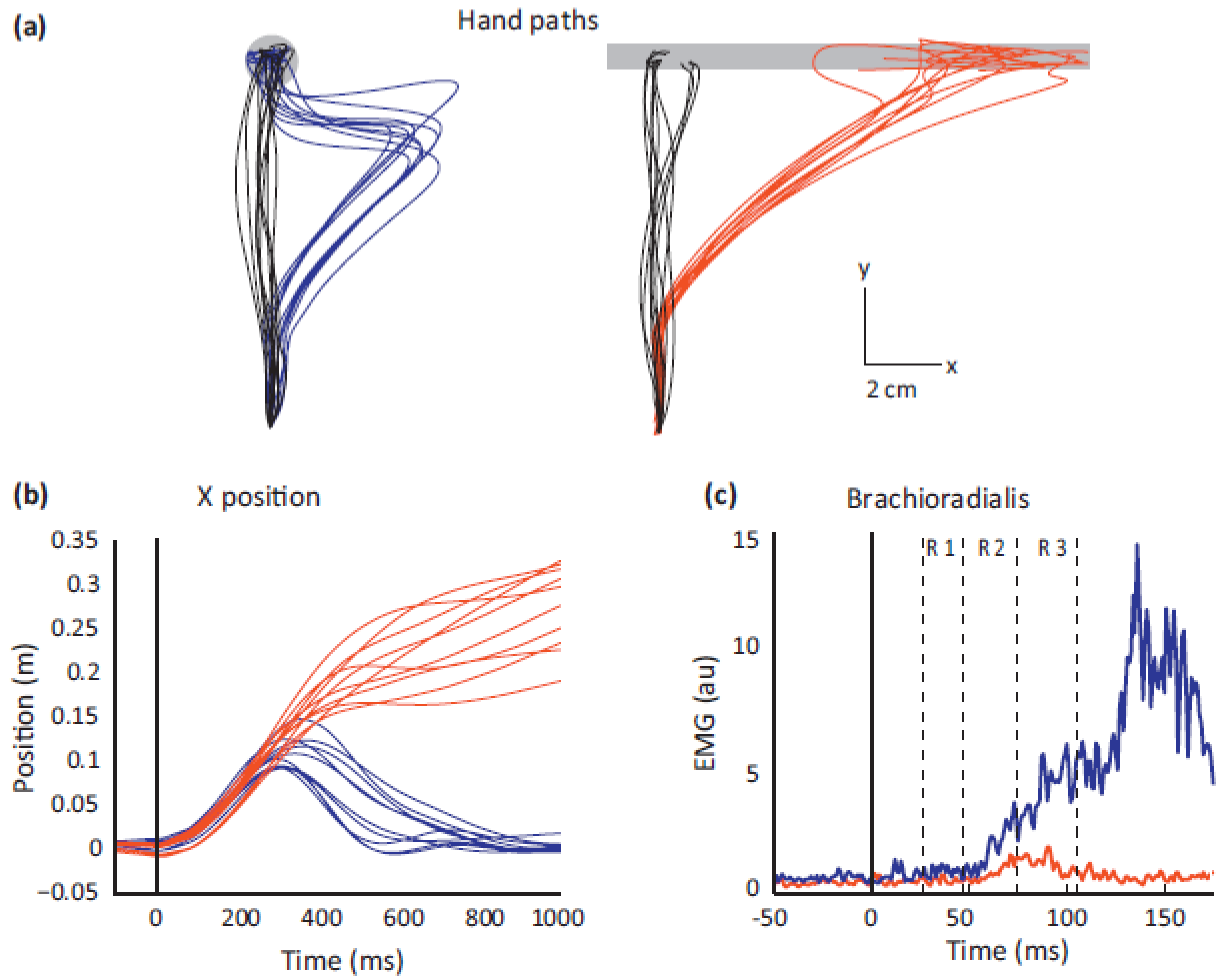

Nashed et al. (2012)

Arm biomechanics affect decisions

- Noise covariances affect movements!

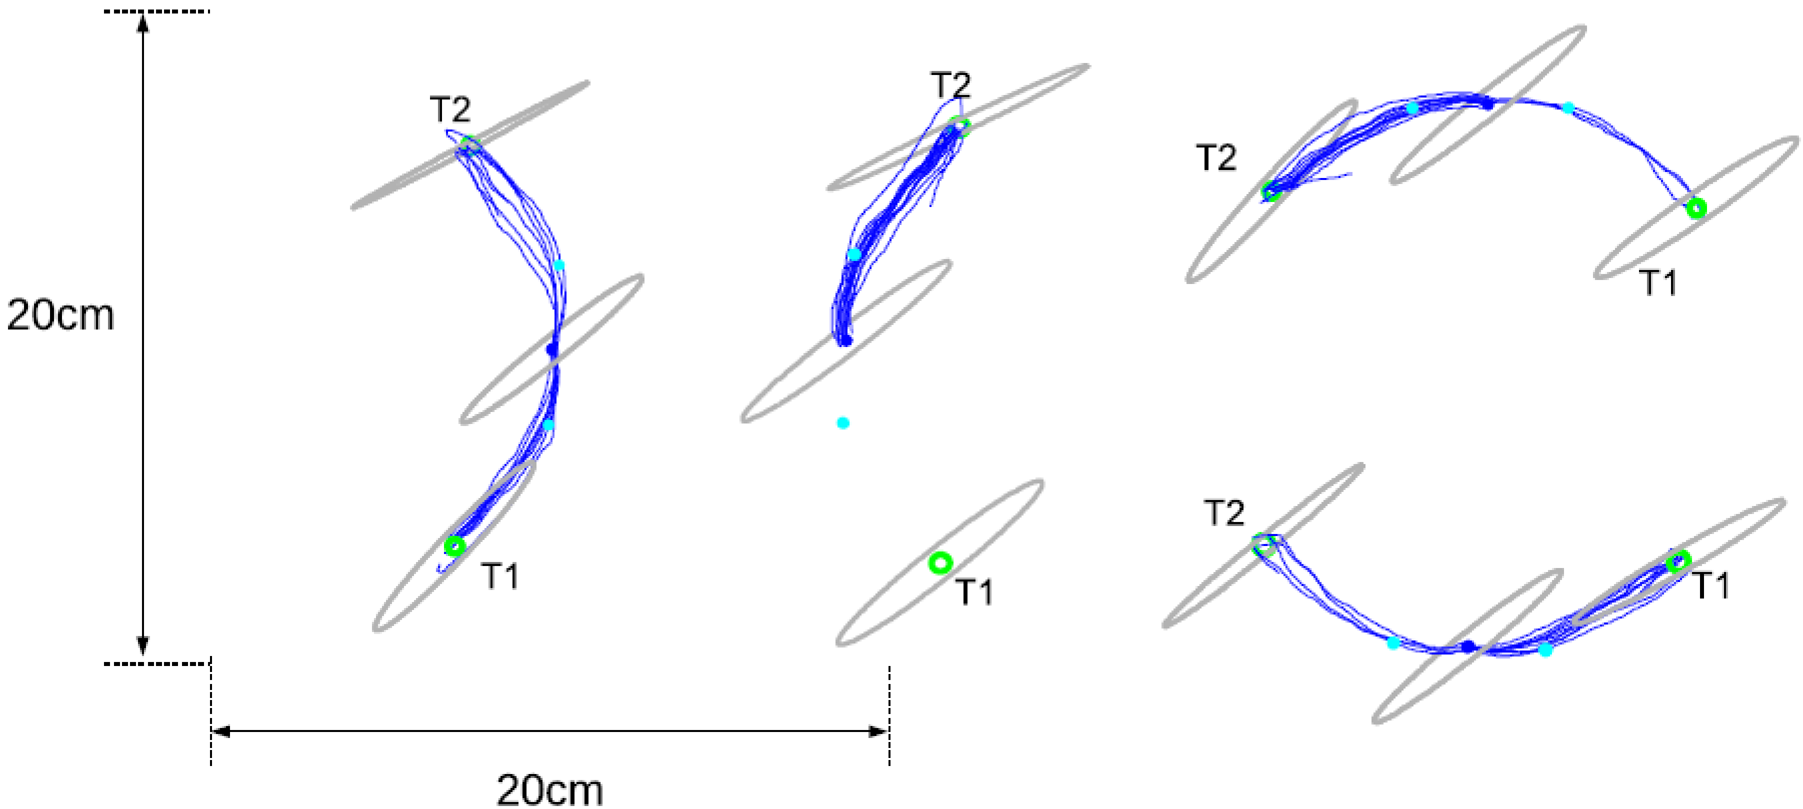

Cos, Belanger, Cisek (2011)

Open discussion

Open discussion

- Interpretation of endpoint errors…
  - Origins from mis-estimations
- Delays
- Infinite horizon control
- Hierarchical control
- Role of biomechanics

Biomechanics

Limb Anatomy

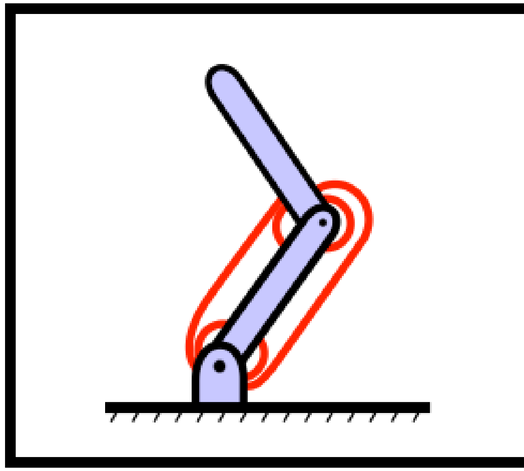

Limb Mechanics

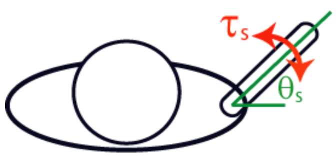

$$\tau_s = I_s \ddot{\theta}_s$$

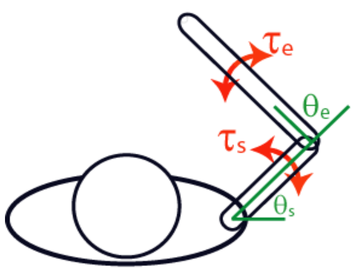

$$\tau_e = \ddot{\theta}_e\left(I_e + \frac{m_e l_e^2}{4}\right) + \ddot{\theta}_s\left(I_e + \frac{m_e l_e^2}{4} + \frac{m_e l_e l_s}{2}\cos(\theta_e)\right) + \left(\frac{m_e l_e l_s}{2}\dot{\theta}_s^2 \sin(\theta_e)\right)$$

$$\tau_s = \ddot{\theta}_s\left(I_s + I_e + m_e l_e l_s \cos(\theta_e) + \frac{m_e l_e^2 + m_s l_s^2}{4} + m_e l_s^2\right) + \ddot{\theta}_e\left(I_e + \frac{m_e l_e^2}{4} + \frac{m_e l_e l_s}{2}\cos(\theta_e)\right) - \left(\frac{m_e l_e l_s}{2}\dot{\theta}_e^2 \sin(\theta_e)\right) - \left(m_e l_e l_s \dot{\theta}_e \dot{\theta}_s \sin(\theta_e)\right)$$

Muscle Mechanics

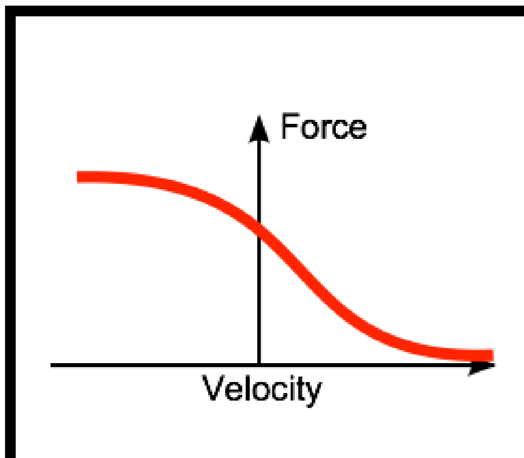

Curtesy of Stephen Scott