

Basics of Control: ICDNS

MSci/MSc

Nick Jones

`nick.jones@imperial.ac.uk`

Review of topics covered in inference course.
How chemical systems can be treated as samplers and can naturally perform probabilistic inference.

Optimal Control of Humanoids

`http://homes.cs.washington.edu/~mordatch/cio/`

This introduction to control is remarkably brief and notably selective. We will cover basic ideas, partly through example, and direct you to proofs in the literature for some points (you should convince yourself of these). Despite this brevity, as before, we will be able to use these concepts to connect to research problems which appear in a biological context.

The next lectures will cover basics of control: feedback, feedforward, stability, controllability, observability, the relevance of linear systems, block diagrams, Bode's integral formula, proportional integral derivative control, internal model control and Kalman filters.

We will then move on to topics in Optimal control.

Control is a point of view

Control is a point of view: it lies within dynamical systems and often involves breaking dynamical systems up into elements like controlling input signals, controlled output signals, feedback, feedforward, and plant (the system being controlled). Because it is a point of view developed in the context of *designed* systems some of the notation and model choices can seem obscure to people who study non-manmade systems. In physical systems we presume that it is the system itself that guides model choice and we are not accustomed to breaking things into parts. For biological systems, though, since we are studying (possibly) *selected* systems, the perspective can be natural, and identifying plant and control is quite sensible. We will also see that the widespread use of linear systems in the theory is very sensible, even if the systems we would like to control are nonlinear.

Because of these issues, make particularly sure you ask me at the end if you don't understand something.

Weak claim. Nature is a designer and so are (control) engineers: we should observe similarities between designed and evolved systems.

More powerful claim. (Optimal) Control can put provable bounds on what constrained systems can achieve. These bounds will limit biological systems and, with the strong assertion that natural systems are optimal, we can expect them to lie on these boundaries. The task then becomes identifying what optimization tasks are relevant.

There is a shortage of brief introductions to control that are both suitable for generic mathematical scientists and impart ready intuition. Most engineering texts assume a skill profile and learning style which isn't a great fit to Mathematicians and Physicists. Ref. [1] is good (free online) but is probably a bit too basic (use of Laplace Transforms and use of the complex plane is only slowly introduced) but very good for examples. Ref. [2] is the best brief introduction I've seen (more for physicists though) but it is dense and requires experience to perceive priority areas. A short introductory account which is certainly worth a look is [3]. I'll be presenting material which draws on both [1, 2, 3]. Note the books by Zabczyk and Sontag which are also good relevant texts [4, 5] (this last is free online and the former free from IC library online).

We will consider dynamical systems $\dot{y} = f(y, u)$ with u a control (or input or strategy) and an output or solution y . We define:

- 1 Open/Closed Loop: $\dot{y} = f(y, u)$ vs $\dot{y} = f(y, g(y))$.
- 2 Feedback: $g(y)$
- 3 Controllability
- 4 (Feedback) Stabilizability
- 5 Observability: $\dot{y} = f(y, u)$ but observe $\nu = h(y)$
- 6 Optimality: often of the form $\int_0^T w(y(t), u(t))dt + W(y(T))$
for observation can, more generally have, $\nu = h(y, u)$.

We imagine a vector of n states y and controls u and n observations ν (in fact all of these three could have different dimensions but we can always pad-out) and consider matrices **A**, **B**, **C**, **D**:

Linear setting

$$\dot{y} = \mathbf{A}y + \mathbf{B}u \quad (1)$$

$$\nu = \mathbf{C}y + \mathbf{D}u \quad (2)$$

we often suppose **D** is zero.

The second equation ν might seem artificial to a theoretical physical scientist - but to an experimentalist it is very natural.

Warmup + the Transfer function

We will consider the Laplace Transform of:

$$\dot{y} = -\omega_0 y(t) + \omega_0 u(t)$$

with $y(0) = 0$. Defining the transfer function:

$$G(s) \equiv \frac{y(s)}{u(s)}$$

in this case $G(s) = \frac{1}{1+s/\omega_0}$. For $s = i\omega$ we will sketch the Bode magnitude plot $|G(i\omega)|$ and phase plot $\arg G(i\omega)$ with ω . We will discuss how the system is acting like a low-pass filter. This kind of view of a linear dynamical system is central for control theory discussions you'll encounter.

A chain of systems

If we feed the output y of our system into another e.g.:

$$\dot{\mu} = -\mu + y(t)$$

It follows that our new output is $\mu(s) = H(s)G(s)u(s)$ i.e. a system with a transfer function $H(s)G(s)$ ($H(s) = \frac{1}{1+s}$ with $\mu(t=0) = 0$). This is a convolution in the time domain.

A block diagram

We'll discuss this notation

$$u(s) \rightarrow \boxed{G(s)} \xrightarrow{y(s)} \boxed{H(s)} \rightarrow \mu(s)$$

One reason why this notation can confuse is that one naively thinks of u merely perturbing the system which has y as its state variable. But here in the Laplace domain $y(s)$ is purely specified by $u(s)$ and the response function.

Feedforward

Control should be simple in principle. If we want our system output to be our control signal then we simply feed our control signal through something that inverts the response function of the system. See the block diagram below: this strategy would involve setting $K(s) = G^{-1}(s)$. Two issues are that we often have disturbances which reflect additional dynamics not accounted for by $G(s)$ and that implementing this inverse might not be possible (e.g. it might be divergent for some s). In practice one thus uses a mix of feedforward and feedback.



Figure : Feedforward

Feedback + block diagram

We suppose we input a control signal $r(t)$ and an error signal $e(t) = r(t) - y(t)$. Our system has dynamics $G(s)$ but we apply a control law $K(s)$ to the error signal. The diagram below is a central object in control theory.




Figure : Feedback

Interpreting the block diagram

From the block diagram we read off $y(s) = K(s)G(s)e(s)$.

Eliminating e we have $y(s) = \frac{K(s)G(s)}{1+K(s)G(s)}r(s) = \frac{L(s)}{1+L(s)}r(s)$.

We can reinterpret $\frac{L(s)}{1+L(s)}$ as the response function of the full dynamical system. $L(s)$ is the response function when there is no control loop.

What is $r(t)$? There is an implicit interpretation about r : it is what we would like our output to be (we want our error $e = r - y$ to be zero). $r(t)$ is thus more than merely an appropriate input to give us our desired output: it *is* our desired output (though this might not be attainable). This helps explain why we would even consider open loop control (and feedforward): if r is something fixed and handed to us (our desired output) we might want to shape it to obtain the right output. If we could pick our input arbitrarily to obtain the right output r then we'd just input $r'(s) = G^{-1}(s)r(s)$ and have no explicit control in our system.

Comparing Open and Closed Loop dynamics

We now investigate the advantage of feedback for a simple example. Fairly obviously, introducing feedback changes our response function from $y/r = L(s)$ to $y/r = \frac{L(s)}{1+L(s)}$.

Let's consider the system/plant from before:

$$\dot{y} = -\omega_0 y(t) + G_0 \omega_0 u(t)$$

with zero i.c.'s so $G(s) = G_0 / (1 + \frac{s}{\omega_0})$. We will now use proportional feedback with the preceding block diagram structure where $u(t) = K_p e(t)$. Thus our dynamics is

$\dot{y} = -\omega_0 y(t) + \omega_0 G_0 K_p (r - y)$. (We'll interpret this).

$y(s)/r(s) = \frac{K_p G_0}{K_p G_0 + 1} \times \frac{1}{1 + \frac{s}{\omega_0(1+K_p G_0)}}$. Compare with $G(s)$ (e.g. we can expect the response frequency to be faster $\omega_0 \rightarrow \omega_0(1 + K_p G_0)$ for appropriately sized G_0).

Which is the control?

We need to be careful about what we call control: I'll labour this since it can confuse. Sometimes the control is called u and we are interested in the open loop response y/u . This makes sense since the thing we manipulate is u . But here, however, we imagine a control signal r and we are interested in y/r the closed loop response to an input signal (we've thus broken down $u = K_p e(t) = K_p(r - y)$ into an internal part $-K_p y$ – the feedback that we can't control once we've set our system up – and the external part r).

Comparing Open and Closed Loop dynamics II

Our open loop dynamics is $\dot{y} = -\omega_0 y(t) + \omega_0 G_0 K_p r$. When $r = \text{constant} = r_c$ then at equilibrium $y = K_p G_0 r$. If we can't guarantee $G_0 = 1/K_p$ (e.g. the gain, G_0 , drifts, as happens in electronics) then we can't ensure that $y \sim r$. By contrast, the equilibrium solution with the closed loop is $y = \frac{K_p G_0}{1 + K_p G_0} r$ which, providing G_0 is sufficiently large, ensures $y \sim r$.

We will see sensitivity appearing later on so let's consider how sensitive the long-time response functions of the open loop ($K_p G_0$) and closed loop ($\frac{K_p G_0}{1 + K_p G_0}$) are to changes in G_0 . Sensitivity of Q w.r.t. P is $S = \frac{P}{Q} \frac{dQ}{dP}$. Convince yourself that for open loop $S = 1$ and for closed $S = \frac{1}{1 + K_p G_0}$. So for $K_p G_0 \gg 1$ the closed loop is very insensitive to gain fluctuations.

This treatment is after [2] and you can explore more there.

Choosing an appropriate controller, $K(s)$, to obtain the desired system response, lies at the heart of control theory.

A block diagram for sources of uncertainty

Here we consider the existence of additive (environmental) disturbances in the output and errors in measurement (sensor). Note that this is only one of a set of possible error models. Note further that our feedback could itself be a dynamical system with some response $H(s)$.




Figure : Feedback with noise sources

It follows that $e = r - (y + n)$; $u = Ke$; $y = Gu + d$

A block diagram for sources of uncertainty II




Figure : Feedback with noise sources

One can find that $y(s) = \frac{KG}{1+KG}(r(s) - n(s)) + \frac{1}{1+KG}d(s)$. The expression $\frac{1}{1+KG} = S(s)$ is a sensitivity: it tells us how much our system responds to disturbances $d(s)$. We will come back to this. While accounting for errors is a central part of control theory we will not investigate this further but point you to robust and adaptive control.

Proportional Integral Derivative (PID) Control

A conventional form of control strategy has the following terms:

$$K(s) = K_p + K_i/s + K_d s$$

We can interpret this as feeding back part of the output error (K_p) something that indicates its rate of change ($K_d s$) and something about how far this error has diverged so far (K_i/s).

We might thus have an ODE like

$$\dot{y} = -\omega_0 y(t) + K_p e + K_i \int_{-\infty}^t e(t') dt' + K_d \frac{de}{dt}.$$

Read the introduction to Sontag [5] and, *in particular*, convince yourself of the role of PID control in stabilizing an inverted pendulum.

PID control will appear in an afternoon session.

Experimentally determining the response function

If I want to control my system, and my control strategy $K(s)$ depends on $G(s)$ (the transfer function of the system before feedback) how do I find $G(s)$? Since $G(s) = y(s)/u(s)$ where y and u are respectively the inputs and outputs of my uncontrolled system I simply concoct a sinusoidal input signal and look at its output. In particular I vary the input frequency and measure the relative amplitude and phase of the output. There are smarter ways, but this works.

Controllability

We call, z , a state, x -reachable if there exists a control u , and time T , such that a system that has output x at $t = 0$ has output z at $t = T$.

If a system is such that for any p, q we find that p is q -reachable then we call the system controllable.

Let's consider a linear dynamical system with a univariate input and output (u and ν scalar) and with feedback $u = -\mathbf{f}^T y$. So $\dot{y} = \mathbf{A}y + u\mathbf{b}$ with $\nu = \mathbf{c}y$ with $\mathbf{f}, \mathbf{b}, \mathbf{c}$ vectors. We can thus define a new dynamical system $\dot{y} = \tilde{\mathbf{A}}y$ s.t. $\tilde{\mathbf{A}} = \mathbf{A} - \mathbf{b}\mathbf{f}^T$.

Controllability for this system is thus about the properties of the matrix $\tilde{\mathbf{A}}$. In particular the kinds of dynamics we can obtain (and whether we can get from p to q for all p, q) depend on the interplay of a constrained choice in \mathbf{f}, \mathbf{b} and the fixed matrix \mathbf{A} . Construct some examples of pairs \mathbf{A}, \mathbf{b} for which there exist pairs p, q where p is not q -reachable.

Controllability and Observability

An example in the case when $x \in \mathbb{R}^2$: \mathbf{A} diagonal (and non-zero) but \mathbf{b} having a zero entry.

For the linear system described to be controllable we require the matrix with columns $\tilde{\mathbf{A}}^i \mathbf{b}$ where $i = 0, \dots, n - 1$ (and is an exponent not an index) with $x \in \mathbb{R}^n$ to be invertible. Look at the treatment of controllability by Zabczyk [3] (or in the other sources provided) and prove that this holds.

We can see that we can have an equivalent problem for whether we can observe the dynamics ν . If, e.g. \mathbf{c} has zero entries. Just as there are general conditions for controllability so too there are conditions for observability (and you'll find discussions of their duality).


We considered the system $\dot{y} = \tilde{\mathbf{A}}y$ s.t. $\tilde{\mathbf{A}} = \mathbf{A} - \mathbf{b}\mathbf{f}^T$. If $\tilde{\mathbf{A}}$ is diagonalizable then we expect the sign of the real part of its eigenvalues to tell us about its stability. Evidently we can partly control the eigenvalues of $\tilde{\mathbf{A}}$ by tuning the feedback \mathbf{f} . Instabilities in systems with feedback can be related to the closed loop transfer function in Laplace space: $y(s)/r(s) = KG/(1 + KG)$. This will diverge when $KG = -1$. A consideration of the poles of transfer functions is thus a canonical topic in control.

This root concept is a substantial part of introductory topics in control but we leave it here.

Feedback linearization

Control problems are frequently posed in terms of linear dynamical systems. This might seem a highly constrained tool-kit. E.g. Why should we care about controllability and observability when they are specified in the linear setting?


It's thus worth knowing about Feedback linearization. This can be understood in terms of two shells of feedback where we have a nonlinear feedback next to the plant that leaves it acting as a linear system that we can control using the tools we've described. The figure below is taken from [1] we'll discuss the slightly different notation you can encounter.



Feedback linearization II

We'll shift notation for this slide to be consistent with the figure (and have a system with state x and observation y - not y and v as before)
We say that a system is feedback linearizable if we can find a control $u = \alpha(x, v) = a(x) + b(x)v$ such that dynamics of the form $\dot{x} = f(x, u) = p(x) + q(x)u$ with $y = h(x)$ becomes linear with input v and output y . E.g. one can sometimes use the form $u = \frac{1}{q(x)}(v - p(x))$.
Feedback linearization is a very rich area and this is an elementary treatment.

For a different approach please read about, and be able to explain, the Ott, Grebogi, and Yorke algorithm for stabilizing chaotic dynamics (you'll find an account of it in ref. [2]).



Internal Model Control

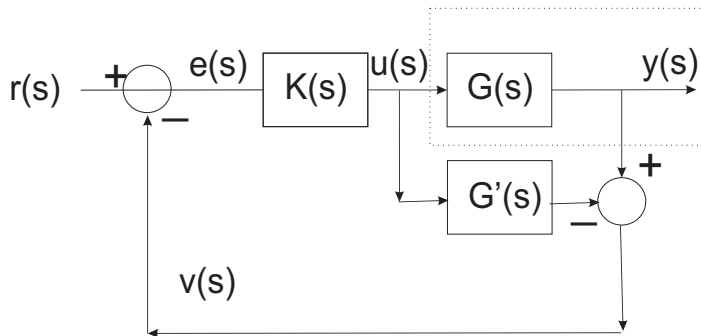



Figure : Internal Model Control (after [2]). Suppose one has a model G' of system G . The strategy is to feedback the difference $((G - G')u)$ between what you suspect the system does to input u ($G'u$) and what it actually does to input u (Gu).

Internal Model Control II

We have response function $\frac{y}{r} = \frac{GK}{1+(G-G')K}$ check this by writing out expressions for e, v, u, y . Which can be reinterpreted as a negative feedback with controller $K'(s) = \frac{1}{1-G'K}$.

Since our feedback is $v = (G - G')y$ we find that a perfect internal model $G' = G$ eliminates the need for feedback, *if* we can invert G' and have feedforward control $K = 1/G'$ (as we've seen before). [If for nonlinear systems we can set $G' \sim G$ then we can hope to linearize]. Internal model control has been taken up as a relevant paradigm in parts of the (motor control) cognitive science community (we will return to this).



Bode's Integral Formula Preliminaries

Also called Bode's sensitivity integral. Related terms are robustness-fragility trade-off or the water-bed effect. We consider the sensitivity of a system with feedback:

$S(s) = \frac{1}{1+L(s)} = \frac{1}{1+K(s)G(s)}$ where $L(s) = K(s)G(s)$ is the open loop gain. S tells us something about the responsiveness of our system at different input frequencies: more specifically it tells us how disturbances $d(s)$ (see the earlier figure with examples of sources of noise d and n) are attenuated. Clearly properties of L put constraints on the sensitivity of the system. We suppose that $L(s) \rightarrow 0$ faster than $\frac{1}{s} \rightarrow 0$ as $s \rightarrow \infty$ and that it has poles p_k in the right half plane.

Bode's Integral Formula

Theorem

Bode's Integral Formula

Given the above constraints on $L(s)$ the following holds:

$$\int_0^{\infty} \log|S(i\omega)|d\omega = \int_0^{\infty} \log\frac{1}{|1 + L(i\omega)|}d\omega = \pi \sum_k p_k$$

Please convince yourself of the proof of this - you'll find it provided in [1].

We will consider the setting when $L(s)$ has no poles in the right-half plane. We can see that whatever K we pick (within this class of L) the total (log) sensitivity must be constant. Thus choosing a $K(s)$ to diminish sensitivity at frequency s' necessitates that there is another frequency s'' for which sensitivity is increased. This is the waterbed effect or robustness-fragility trade off. This conservation principle is generic and can be expected to constrain biological systems [6].

Linear setting

Dynamics with additive noise in our output d :

$$y_{k+1} = \phi y_k + u_k + d_k$$

Observation with additive noise in our observation n :

$$\nu_{k+1} = y_{k+1} + n_{k+1}$$

With our state variable y , control u and observation data ν (and time index k).

Discrete Time Control with Noise II

$$y_{k+1} = \phi y_k + u_k + d_k$$

$$\nu_{k+1} = y_{k+1} + n_{k+1}$$

If we have a conventional negative feedback setup then the signal we feedback is our current observation $\nu = y + n$. This presents a problem since our error should be $e = r - y$ but actually it's $e = r - (y + n)$. This means that a control system that tries to set $e = 0$ will not set $y = r$: it will be very susceptible to non-zero n .

Discrete Time Control with Noise III

$$y_{k+1} = \phi y_k + u_k + d_k$$

$$\nu_{k+1} = y_{k+1} + n_{k+1}$$

We will see that the Kalman filter helps us be smart about errors in our observations in order to correctly work out the state y . It combines our current state observations ν_{k+1} with what you would have predicted \hat{y}_{k+1} about what the state actually is (y_{k+1}). That means that we can limit pollution by n_k .

This is the heart of Kalman filters: *be smart, don't just use your current (noise polluted) observation ν as your best guess as to y , instead combine it with some past information that might have averaged out n . How much you weight your current predicted state vs your latest observation depends on how much you trust your current information.*

Since it combines estimation with feedback the Kalman filter lies at the interface of control and inference.

Introducing the (Univariate) Kalman filter

Kalman Filter

Actual state: $y_{k+1} = \phi y_k + u_k + d_k$

Actual Observation: $\nu_{k+1} = y_{k+1} + n_{k+1}$

Predicted state (which evolves your last best estimate in time):

$$\hat{y}_{k+1} = \phi \hat{y}_{k|k} + u_k$$

Current Best estimate (combine predicted state with current observation):

$$\hat{y}_{k+1|k+1} \equiv (1 - K) \hat{y}_{k+1} + K \nu_{k+1}$$

Predicted observation: $\hat{\nu}_{k+1} = \hat{y}_{k+1}$

I'm following [2] in this account: you'll find a closely complementary variant there (but watch out for shifted y 's, n 's and x 's).

Introducing the (Univariate) Kalman filter II

$$\text{State : } y_{k+1} = \phi y_k + u_k + d_k$$

$$\text{Observe : } \nu_{k+1} = y_{k+1} + n_{k+1}$$

$$\text{Predict : } \hat{y}_{k+1} = \phi \hat{y}_{k|k} + u_k$$

$$\text{Estimate : } \hat{y}_{k+1|k+1} \equiv (1 - K) \hat{y}_{k+1} + K \nu_{k+1}$$

$$\text{Predicted Observation : } \hat{\nu}_{k+1} = \hat{y}_{k+1}$$

Questions

- So what are we filtering? We are attempting to improve our estimate of y and so filter out observation error n . We often use this improved estimate $\hat{y}_{k+1|k+1}$ to feedback into our control ensuring that $e = r - \hat{y}_{k+1|k+1} \sim r - y$.
- When is this inference architecture used? d and n are assumed Gaussian distributed and memoryless stochastic processes (white noise).

Introducing the (Univariate) Kalman filter III

State : $y_{k+1} = \phi y_k + u_k + d_k$

Observe : $\nu_{k+1} = y_{k+1} + n_{k+1}$

Predict : $\hat{y}_{k+1} = \phi \hat{y}_{k|k} + u_k$

Estimate : $\hat{y}_{k+1|k+1} \equiv (1 - K)\hat{y}_{k+1} + K\nu_{k+1}$

Predicted Observation : $\hat{\nu}_{k+1} = \hat{y}_{k+1}$

Questions

- K ? This is the trade-off constant between our predictions for the current state y based on past evidence (\hat{y}_{k+1}) and our latest data ν_{k+1} . Our best estimate of the current state (defined as $\hat{y}_{k+1|k+1}$). How to pick K ? See next slide.
- What is \hat{y}_{k+1} ? This is our prediction as to the current state if we believe our best estimate $\hat{y}_{k|k}$ from last time and assume a noiseless evolution according to the control u_k we have applied. We assume we know u_k (and ϕ).

Kalman Gain

We can define our error in estimating the state as $er_k = y_k - \hat{y}_{k|k}$ (this is unknown to us since we never see y_k).

We want to select our Kalman gain K so as to have our error as tightly distributed about zero as possible i.e. to minimize $\langle e^2 \rangle$. By first showing that $er_{k+1} = (1 - K)[\phi er_k + d_k] - Kn_{k+1}$ find an expression in $\langle er_{k+1}^2 \rangle$ and $\langle er_k^2 \rangle$ (using the independence of the processes er, d, n) which can be minimized (w.r.t. K) to show that:

$$K = \frac{\phi^2 \langle er_k^2 \rangle + \bar{d}^2}{\phi^2 \langle er_k^2 \rangle + \bar{d}^2 + \bar{n}^2}$$

Where \bar{d} and \bar{n} are expectations of the corresponding processes.

Recall that $\hat{y}_{k+1|k+1} \equiv (1 - K)\hat{y}_{k+1} + Kn_{k+1}$.

If environmental noise dominates then $\bar{d} \gg \bar{n}$. This means $K \sim 1$ and suggests we can forget the filter and trust that ν is close to y . By contrast, if $\bar{n} \gg \bar{d}$ then we really need the filter, since our estimates are relatively very untrustworthy.

Generalizations (Kalman Filter)

You'll find that the Kalman filter generalizes straightforwardly to the multivariate case.

We note that if we take our timesteps sufficiently small then nonlinear dynamics can still be incorporated within the linear framework of the Kalman filter (Taylor expand). This is called the extended Kalman filter.

Topic Review + Practical

I'll run through a reminder of the topics we've covered.

Implement a univariate Kalman filter. For given noise processes n and d find the squared error in the filter as we vary the Kalman gain.

Bibliography

- [1] K.J. Åström and R.J. Murray, Feedback systems, Princeton University Press. 2008. Free online.
- [2] J. Bechhoefer, Feedback for physicists: A tutorial essay on control, Reviews of Modern Physics, 77, 783, 2005. Free to you online.
- [3] J. Zabczyk, Classical control theory, 2001. Lecture notes online.
- [4] J. Zabczyk, Mathematical control theory, Birkhauser, 1995. Free to IC students.
- [5] E. Sontag, Mathematical control theory, Springer, 1998. Free online.
- [6] M. Csete and J. Doyle, Reverse engineering of biological complexity, Science 295, 1664, 2002.