Lecture Notes 7:
Difference and Differential Equations

Peter J. Hammond

Autumn 2012; revised 2013 and 2014

Lecture Outline

Introduction

Difference vs. Differential Equations

First-Order Difference Equations

Systems of Linear Difference Equations

Diagonalizing a Non-Symmetric Matrix

Outline

Walking as a Simple Difference Equation

What is the difference
between difference and differential equations?

Walking on two feet is a discrete time process,
with time domain $T = \{0, 1, 2, \ldots\} = \mathbb{Z}_+ = \{0\} \cup \mathbb{N}$.

After m steps, the respective positions $\ell, r \in \mathbb{R}^2$
of the left and right feet on the ground
can be described by the two functions $T \ni m \mapsto (\ell_m, r_m)$.

It is impossible for a step to be longer than one's stride length.

So the walking process that starts with the left foot
might be described by the two coupled equations

$$\ell_m = \begin{cases} \lambda(r_{m-1}) & \text{if } m \text{ is odd} \\ \ell_{m-1} & \text{if } m \text{ is even} \end{cases} \quad \text{and} \quad r_m = \begin{cases} \rho(\ell_{m-1}) & \text{if } m \text{ is even} \\ r_{m-1} & \text{if } m \text{ is odd} \end{cases}$$

for $m = 0, 1, 2, \ldots$.

Walking as a More Complicated Difference Equation

Or, if the direction of each pace
is affected by the direction of its predecessor, by

$$\ell_m = \begin{cases} \lambda(r_{m-1}, \ell_{m-1}) & \text{if } m \text{ is odd} \\ \ell_{m-1} & \text{if } m \text{ is even} \end{cases}$$

$$\text{and} \quad r_m = \begin{cases} \rho(\ell_{m-1}, r_{m-1}) & \text{if } m \text{ is even} \\ r_{m-1} & \text{if } m \text{ is odd} \end{cases}$$

for $m = 0, 1, 2, \ldots$.

Walking as a Differential Equation

By constrast, a walker's centre of mass
must be a continuous function of time,
described by a mapping $\mathbb{R}_+ \ni t \mapsto (x(t), y(t), z(t)) \in \mathbb{R}^3$.

The time domain is therefore $T := \mathbb{R}_+$.

The same will be true for the position of, for instance,
the walker's left big toenail.

Actually, the motion could be lethal
unless the function $\mathbb{R}_+ \ni t \mapsto (x(t), y(t), z(t)) \in \mathbb{R}^3$
has a continuous third derivative $\dfrac{d^3}{dt^3}(x(t), y(t), z(t))$.

It is relatively common to indicate by:

- **a subscript** a discrete time function like $m \mapsto x_m$;
- **parentheses** a continuous time function like $t \mapsto x(t)$.

First-Order Difference Equations

Let $T \ni t \mapsto x_t \in X$ describe a discrete time process, with $X = \mathbb{R}$ (or $X = \mathbb{R}^m$) as the state space.

Its difference at time t is defined as

$$\Delta x_t := x_{t+1} - x_t$$

A standard first-order difference equation takes the form

$$x_{t+1} - x_t = \Delta x_t = d_t(x_t)$$

where each $d_t : X \to X$, or equivalently,

$$T \times X \ni (t, x) \mapsto d_t(x)$$

Obviously, it is equivalent to the recurrence relation

$$T \times X \ni (t, x) \mapsto r_t(x)$$

where $r_t(x) = x + d_t(x)$, or equivalently, $d_t(x) = r_t(x) - x$.

Equivalent Recurrence Relations

Thus difference equations and recurrence relations are entirely equivalent.

We follow standard mathematical practice in using the notation for recurrence relations, even when discussing difference equations.

Linear Equations

A simple linear equation for a finite time horizon T takes the form

$$x_{t+1} - x_t = d_t \qquad (t = 0, 1, \ldots, T-1)$$

where the differences are constants $d_t \in \mathbb{R}$, independent of x.

Provided that $T \geq 6$, the matrix form of this equation is

$$\begin{pmatrix} -1 & 1 & 0 & \ldots & 0 & 0 & 0 \\ 0 & -1 & 1 & \ldots & 0 & 0 & 0 \\ 0 & 0 & -1 & \ldots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & -1 & 1 & 0 \\ 0 & 0 & 0 & \ldots & 0 & -1 & 1 \end{pmatrix} \begin{pmatrix} x_0 \\ x_1 \\ x_2 \\ \vdots \\ x_{T-2} \\ x_{T-1} \\ x_T \end{pmatrix} = \begin{pmatrix} d_0 \\ d_1 \\ d_2 \\ \vdots \\ d_{T-2} \\ d_{T-1} \end{pmatrix}$$

Matrix Form

The matrix form of the difference equation is $\mathbf{Dx} = \mathbf{d}$, where:

1. $\mathbf{D}$ is the $T \times (T+1)$ difference matrix whose coefficients are

$$d_{st} = \begin{cases} -1 & \text{if } t = s \\ 1 & \text{if } t = s+1 \\ 0 & \text{otherwise} \end{cases}$$

   for $s = 0, 1, 2, \ldots, T$ and $t = 1, 2, \ldots, T$;
2. $\mathbf{x}$ is the $T+1$-dimensional column vector $(x_t)_{t=0}^{T}$ of endogenous unknowns, to be determined;
3. $\mathbf{d}$ is the T-dimensional column vector $(d_t)_{t=0}^{T-1}$ of exogenous shocks.

Partitioned Matrix Form

The matrix equation $\mathbf{Dx} = \mathbf{d}$ can be written in partitioned form as

$$\begin{pmatrix} \mathbf{U} & \mathbf{e}_T \end{pmatrix} \begin{pmatrix} \mathbf{x}^{T-1} \\ x_T \end{pmatrix} = \mathbf{d}$$

where:

1. $\mathbf{U}$ is an upper triangular $T \times T$ matrix;
2. $\mathbf{e}_T = (0, 0, 0, \ldots, 0, 1)$ is the Tth unit vector of the space $\mathbb{R}^T$;
3. $\mathbf{x}^{T-1}$ denotes the column vector whose transpose $(\mathbf{x}^{T-1})^\top$ equals the row T-vector $(x_0, x_1, x_2, \ldots, x_{T-2}, x_{T-1})$.

In fact the matrix $-\mathbf{U}$ is even upper unitriangular.

Hence there are T independent equations in $T + 1$ unknowns, leaving one degree of freedom in the solution.

An Initial Condition

Consider the difference equation $x_{t+1} - x_t = d_t$,
or $\mathbf{Dx} = \mathbf{d}$ in matrix form.

An initial condition specifies an exogenous value $\bar{x}_0$
for the value x_0 at time 0.

This removes the only degree of freedom
in the system of T equations in $T+1$ unknowns.

The obvious unique solution is then that each x_t is the forward sum

$$x_t = \bar{x}_0 + \sum_{s=0}^{t-1} d_s$$

of the initial state $\bar{x}_0$, and of the t exogenously specified
succeeding differences d_s $(s = 0, 1, \ldots, t-1)$.

Terminal Condition

Alternatively, a terminal condition
for the difference equation $x_{t+1} - x_t = d_t$
specifies an exogenous value $\bar{x}_T$ for the value x_T at time T.

It leads to a unique solution as a backward sum

$$x_t = \bar{x}_T - \sum_{s=0}^{T-t-1} d_{T-s}$$

of the initial state $\bar{x}_T$, and of the $T - t$ exogenously specified
preceding backward differences $-d_{T-s}$ $(s = 0, 1, \ldots, T - t - 1)$.

Particular and General Solutions

We are interested in solving the system $\mathbf{D}\mathbf{x} = \mathbf{d}$
of T equations in $T + 1$ unknowns,
where $\mathbf{D}$ is a $T \times (T + 1)$ matrix.

When the rank of $\mathbf{D}$ is T, there is one degree of freedom.

The homogeneous equation $\mathbf{D}\mathbf{x} = \mathbf{0}$
will have a one-dimensional space of solutions $x_t^H = \xi \bar{x}_t^H$ $(\xi \in \mathbb{R})$.

Given any particular solution x_t^P satisfying $\mathbf{D}\mathbf{x}^P = \mathbf{d}$
for the particular time series $\mathbf{d}$ of forcing terms,
the general solution x_t^G must also satisfy $\mathbf{D}\mathbf{x}^G = \mathbf{d}$.

Simple subtraction leads to $\mathbf{D}(\mathbf{x}^G - \mathbf{x}^P) = \mathbf{0}$, so $\mathbf{x}^G - \mathbf{x}^P = \mathbf{x}^H$
for some solution $\mathbf{x}^H$ of the homogeneous equation $\mathbf{D}\mathbf{x} = \mathbf{0}$.

So $\mathbf{x}$ solves the equation $\mathbf{D}\mathbf{x} = \mathbf{d}$
iff there exists a scalar $\xi \in \mathbb{R}$ such that $\mathbf{x} = \mathbf{x}^P + \xi\, \mathbf{x}^H$,
which leads to the formula $\mathbf{x}^G = \mathbf{x}^P + \xi\, \mathbf{x}^H$ for the general solution.

Application: Wealth Accumulation in Discrete Time

Consider a consumer who, in discrete time $t = 0, 1, 2, \ldots$:

- starts each period t
  with an amount w_t of accumulated wealth;
- receives income y_t;
- spends an amount e_t;
- earns interest on the residual wealth $w_t + y_t - e_t$ at the rate r_t.

The process of wealth accumulation is then described
by any of the equivalent equations

$$w_{t+1} = (1 + r_t)(w_t + y_t - e_t) = \rho_t(w_t - x_t) = \rho_t(w_t + s_t)$$

where, at each time t,

- $\rho_t := 1 + r_t$ is the interest factor;
- $x_t = e_t - y_t$ denotes net expenditure;
- $s_t = y_t - e_t = -x_t$ denotes net saving.

Compound Interest

Define the compound interest factor

$$R_t := \prod_{k=0}^{t-1}(1 + r_k) = \prod_{k=0}^{t-1} \rho_k$$

with the convention that the product of zero terms equals 1.

It is the unique solution
to the recurrence relation $R_{t+1} = (1 + r_t)R_t$
that satisfies the initial condition $R_0 = 1$.

In the special case when $r_t = r$ (all t),
it reduces to $R_t = (1 + r)^t = \rho^t$.

Present Discounted Value (PDV)

We transform the difference equation $w_{t+1} = \rho_t(w_t - x_t)$
by using the compound interest factor $R_t = \prod_{k=0}^{t-1} \rho_k$
in order to discount both future wealth and expenditure.

To do so, define new variables ω_t, ξ_t
for the present discounted values (PDVs) of, respectively:

1. wealth w_t at time t as $\omega_t := (1/R_t)w_t$;
2. net expenditure x_t at time t as $\xi_t := (1/R_t)x_t$.

With these new variables, the wealth equation $w_{t+1} = \rho_t(w_t - x_t)$
becomes

$$R_{t+1}\omega_{t+1} = \rho_t R_t(\omega_t - \xi_t)$$

But $R_{t+1} = \rho_t R_t$, so eliminating this common factor
reduces the equation to $\omega_{t+1} = \omega_t - \xi_t$,
with the evident solution $\omega_t = \omega_0 - \sum_{k=0}^{t-1} \xi_k$ for $k = 1, 2, \ldots$.

Systems of Linear Difference Equations

Most economic models, especially econometric models, involve simultaneous time series for several different variables.

Consider a first-order linear difference equation

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

for an n-dimensional process $T \ni t \mapsto \mathbf{x}_t \in \mathbb{R}^n$,
where each matrix $\mathbf{A}_t$ is $n \times n$.

We will prove by induction on t that for $t = 0, 1, 2, \ldots$
there exist suitable matrices $\mathbf{P}_{t,k}$ $(k = 0, 1, 2, \ldots, t)$ such that,
given any possible value of the initial state vector $\mathbf{x}_0$
and of the forcing terms $\mathbf{d}_t$ $(t = 0, 1, 2, \ldots)$,
the unique solution can be expressed as

$$\mathbf{x}_t = \mathbf{P}_{t,0}\mathbf{x}_0 + \sum\nolimits_{k=1}^{t} \mathbf{P}_{t,k}\mathbf{d}_{k-1}$$

The proof, of course,
will also involve deriving a recurrence relation for these matrices.

Early Terms of the Matrix Solution

Because $\mathbf{x}_0 = \mathbf{P}_{0,0}\mathbf{x}_0 = \mathbf{x}_0$,
the first term is obviously $\mathbf{P}_{0,0} = \mathbf{I}$ when $t = 0$.

Next $\mathbf{x}_1 = \mathbf{A}_0\mathbf{x}_0 + \mathbf{d}_0$ when $t = 1$
implies that $\mathbf{P}_{1,0} = \mathbf{A}_0$, $\mathbf{P}_{1,1} = \mathbf{I}$.

Next, the solution for $t = 2$ is

$$\mathbf{x}_2 = \mathbf{A}_1\mathbf{x}_1 + \mathbf{d}_1 = \mathbf{A}_1\mathbf{A}_0\mathbf{x}_0 + \mathbf{A}_1\mathbf{d}_0 + \mathbf{d}_1$$

This formula matches the formula

$$\mathbf{x}_t = \mathbf{P}_{t,0}\mathbf{x}_0 + \sum\nolimits_{k=1}^{t} \mathbf{P}_{t,k}\mathbf{d}_{k-1}$$

when $t = 2$ provided that:

- $\mathbf{P}_{2,0} = \mathbf{A}_1\mathbf{A}_0$;
- $\mathbf{P}_{2,1} = \mathbf{A}_1$;
- $\mathbf{P}_{2,2} = \mathbf{I}$.

Matrix Solution

Now, substituting the two expansions

$$\begin{array}{rrcl} & \mathbf{x}_t & = & \mathbf{P}_{t,0}\mathbf{x}_0 + \sum_{k=1}^{t} \mathbf{P}_{t,k}\mathbf{d}_{k-1} \\ \text{and} & \mathbf{x}_{t+1} & = & \mathbf{P}_{t+1,0}\mathbf{x}_0 + \sum_{k=1}^{t+1} \mathbf{P}_{t+1,k}\mathbf{d}_{k-1} \end{array}$$

into both sides of the original equation $\mathbf{x}_{t+1} = \mathbf{A}_t\mathbf{x}_t + \mathbf{d}_t$ gives

$$\mathbf{P}_{t+1,0}\mathbf{x}_0 + \sum_{k=1}^{t+1} \mathbf{P}_{t+1,k}\mathbf{d}_{k-1} = \mathbf{A}_t \left(\mathbf{P}_{t,0}\mathbf{x}_0 + \sum_{k=1}^{t} \mathbf{P}_{t,k}\mathbf{d}_{k-1} \right) + \mathbf{d}_t$$

Equating the matrix coefficients of $\mathbf{x}_0$ and of each $\mathbf{d}_{k-1}$ implies that for general t one has $\mathbf{P}_{t+1,k} = \mathbf{A}_t\mathbf{P}_{t,k}$ for $k = 0, 1, \ldots, t+1$.

This equation implies that

$$\begin{aligned} \mathbf{P}_{t,\,0} &= \mathbf{A}_{t-1} \cdot \mathbf{A}_{t-2} \cdots \mathbf{A}_0 \\ \mathbf{P}_{t,\,k} &= \mathbf{A}_{t-1} \cdot \mathbf{A}_{t-2} \cdots \mathbf{A}_k \\ \mathbf{P}_{t,\,t} &= \mathbf{I} \end{aligned}$$

or, after defining the product of the empty set of matrices as $\mathbf{I}$,

$$\mathbf{P}_{t,\,k} = \prod_{s=1}^{t-k} \mathbf{A}_{t-s}$$

Constant Coefficients

In the case of constant coefficients,
the products reduce to powers.

Specifically, $\mathbf{P}_{t,\,k} = \mathbf{A}^{t-k}$, where $\mathbf{A}^0 = \mathbf{I}$.

The solution to $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{d}_t$ is therefore

$$\mathbf{x}_t = \mathbf{A}^t \mathbf{x}_0 + \sum_{k=1}^{t} \mathbf{A}^{t-k} \mathbf{d}_k$$

The Autonomous Case

The general first-order equation in $\mathbb{R}^n$
can be written as $\mathbf{x}_{t+1} = \mathbf{F}_t(\mathbf{x}_t)$
where $T \times \mathbb{R}^n \ni (t, \mathbf{x}) \mapsto \mathbf{F}_t(\mathbf{x}) \in \mathbb{R}^n$.

In the autonomous case, the function $(t, \mathbf{x}) \mapsto \mathbf{F}_t(\mathbf{x})$
reduces to $\mathbf{x} \mapsto \mathbf{F}(\mathbf{x})$, independent of t.

In the linear case with constant coefficients,
the function $\mathbf{x} \mapsto \mathbf{F}(\mathbf{x})$ takes the affine form $\mathbf{F}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{d}$.

That is, $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{d}$.

Linear Case with Constant Coefficients

Given the equation $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{d}$,
the earlier formula for the solution leads to

$$\mathbf{x}_t = \mathbf{A}^t \mathbf{x}_0 + \sum_{k=1}^{t} \mathbf{A}^{t-k} \mathbf{d} = \mathbf{A}^t \mathbf{x}_0 + \mathbf{S}_t \mathbf{d}$$

where the matrix

$$\mathbf{S}_t := \mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \cdots + \mathbf{A}^{t-1} = \sum_{k=1}^{t} \mathbf{A}^{t-k}$$

is the matrix analogue of the geometric series

$$\begin{aligned} s_t &:= 1 + a + a^2 + \cdots + a^{t-1} \\ &= \sum_{k=1}^{t} a^{t-k} = \begin{cases} \dfrac{1 - a^t}{1 - a} & \text{if } a \neq 1 \\ t & \text{if } a = 1 \end{cases} \end{aligned}$$

Summing the Geometric Series

Recall the trick for finding $s_t := 1 + a + a^2 + \cdots + a^{t-1}$ is to multiply each side by $1 - a$.

Because all terms except the first and last cancel, this shows that $(1 - a)s_t = 1 - a^t$ and so $s_t = (1 - a)^{-1}(1 - a^t)$ provided that $a \neq 1$.

Applying the same trick to $\mathbf{S}_t := \mathbf{I} + \mathbf{A} + \mathbf{A}^2 + \cdots + \mathbf{A}^{t-1}$ yields $(\mathbf{I} - \mathbf{A})\mathbf{S}_t = \mathbf{I} - \mathbf{A}^t$.

Provided that $(\mathbf{I} - \mathbf{A})^{-1}$ exists, we can pre-multiply each side by this inverse to get $\mathbf{S}_t = (\mathbf{I} - \mathbf{A})^{-1}(\mathbf{I} - \mathbf{A}^t)$.

This leads to the solution

$$\mathbf{x}_t = \mathbf{A}^t\mathbf{x}_0 + \mathbf{S}_t\mathbf{d} = \mathbf{A}^t\mathbf{x}_0 + (\mathbf{I} - \mathbf{A})^{-1}(\mathbf{I} - \mathbf{A}^t)\mathbf{d}$$

Stationary States

Given an autonomous equation $\mathbf{x}_{t+1} = \mathbf{F}(\mathbf{x}_t)$,
a stationary state is a fixed point $\mathbf{x}^* \in \mathbb{R}^n$ of the mapping $\mathbf{F}$.

It earns its name because if $\mathbf{x}_s = \mathbf{x}^*$ for any finite s,
then $\mathbf{x}_t = \mathbf{x}^*$ for all $t = s, s+1, \ldots$.

Wherever it exists, the solution of the autonomous equation can be written as a function $\mathbf{x}_t = \Phi_{t-s}(\mathbf{x}_s)$ $(t = s, s+1, \ldots)$ of the state $\mathbf{x}_s$ at time s,
as well as of the number of periods $t - s$ that the function $\mathbf{F}$ must be iterated in order to determine the state $\mathbf{x}_t$ at time t.

Indeed, the sequence of functions $\Phi_k : \mathbb{R}^n \to \mathbb{R}^n$ $(k \in \mathbb{N})$
is defined iteratively by $\Phi_k(\mathbf{x}) = \mathbf{F}(\Phi_{k-1}(\mathbf{x}))$ for all $\mathbf{x}$.

Note that any stationary state $\mathbf{x}^*$ is a fixed point
of each mapping Φ_k in the sequence, as well as $\Phi_1 \equiv \mathbf{F}$.

Local and Global Stability

The stationary state $\mathbf{x}^*$ is:

- globally stable if $\Phi_k(\mathbf{x}_0) \to \mathbf{x}^*$ as $k \to \infty$, regardless of the initial state $\mathbf{x}_0$;
- locally stable if there is an (open) neighbourhood $N \subset \mathbb{R}^n$ of $\mathbf{x}^*$ such that whenever $\mathbf{x}_0 \in N$ one has $\Phi_k(\mathbf{x}_0) \to \mathbf{x}^*$ as $k \to \infty$.

We begin by studying linear systems,
for which local stability is equivalent to global stability.

Later, we will consider the local stability of non-linear systems.

Stability in the Linear Case

Recall that the autonomous linear equation
takes the form $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{d}$.

The vector $\mathbf{x}^* \in \mathbb{R}^n$ is a stationary state
if and only if $\mathbf{x}_t = \mathbf{x}^* \implies \mathbf{x}_{t+1} = \mathbf{x}^*$,
which is true if and only if $\mathbf{x}^* = \mathbf{A}\mathbf{x}^* + \mathbf{d}$,
or iff $\mathbf{x}^*$ solves the linear equation $(\mathbf{I} - \mathbf{A})\mathbf{x} = \mathbf{d}$.

Of course, if the matrix $\mathbf{I} - \mathbf{A}$ is singular,
then there could either be no stationary state,
or a continuum of stationary states.

For simplicity, we assume that $\mathbf{I} - \mathbf{A}$ has an inverse.

Then there is a unique stationary state $\mathbf{x}^* = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{d}$.

Homogenizing the Linear Equation

Given the equation $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{d}$
and the stationary state $\mathbf{x}^* = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{d}$,
define the new state as the deviation $\mathbf{y} := \mathbf{x} - \mathbf{x}^*$
of the state $\mathbf{x}$ from the stationary state $\mathbf{x}$.

This transforms the original equation $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{d}$ to

$$\mathbf{y}_{t+1} + \mathbf{x}^* = \mathbf{A}(\mathbf{y}_t + \mathbf{x}^*) + \mathbf{d} = \mathbf{A}\mathbf{y}_t + \mathbf{A}\mathbf{x}^* + \mathbf{d}$$

Because the stationary state satisfies $\mathbf{x}^* = \mathbf{A}\mathbf{x}^* + \mathbf{d}$,
this reduces the original equation $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{d}$
to the homogeneous equation $\mathbf{y}_{t+1} = \mathbf{A}\mathbf{y}_t$,
whose obvious solution is $\mathbf{y}_t = \mathbf{A}^t\mathbf{y}_0$.

Stability in the Diagonal Case

Suppose that $\mathbf{A}$ is the diagonal matrix $\mathbf{\Lambda} = \mathbf{diag}(\lambda_1, \lambda_2, \ldots, \lambda_n)$.

Then the powers are easy:

$$\mathbf{A}^t = \mathbf{\Lambda}^t = \mathbf{diag}(\lambda_1^t, \lambda_2^t, \ldots, \lambda_n^t)$$

The "homogenized" vector equation $\mathbf{y}_t = \mathbf{A}\mathbf{y}_{t-1}$
can be expressed component by component as the set

$$y_{i,t} = \lambda_i y_{i,t-1} \quad (i = 1, 2, \ldots, n)$$

of n uncoupled difference equations in one variable.

The solution of $\mathbf{y}_t = \mathbf{A}\mathbf{y}_{t-1}$ with $\mathbf{y}_0 = \mathbf{z} = (z_1, z_2, \ldots, z_n)$
is then $\mathbf{y}_t = (\lambda_1^t z_1, \lambda_2^t z_2, \ldots, \lambda_n^t z_n)$.

Hence $\mathbf{y}_t \to \mathbf{0}$ holds for all $\mathbf{y}_0$
if and only if $|\lambda_i| < 1$ for $i = 1, 2, \ldots, n$.

Recall that when $\lambda = \alpha \pm i\beta$, one has $|\lambda| = \sqrt{\alpha^2 + \beta^2}$.

Warning Example

Consider the 2×2 matrix $\mathbf{A} = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 2 \end{pmatrix}$.

The solution of the difference equation $\mathbf{y}_t = \mathbf{A}\mathbf{y}_{t-1}$ with $\mathbf{y}_0 = \mathbf{z} = (z_1, z_2, \ldots, z_n)$ is then

$$\mathbf{y}_t = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 2 \end{pmatrix}^t \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \begin{pmatrix} 2^{-t} & 0 \\ 0 & 2^t \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \begin{pmatrix} 2^{-t} z_1 \\ 2^t z_2 \end{pmatrix}$$

Then $\mathbf{y}_t \to 0$ as $t \to \infty$ provided that $z_2 = 0$.

But the norm $\|\mathbf{y}_t\| \to +\infty$ whenever $z_2 \neq 0$.

In this case one says that $\mathbf{A}$ exhibits saddle point stability because starting with $z_2 = 0$ allows convergence, but starting with $z_2 \neq 0$ ensures divergence.

This explains why one says that the $n \times n$ matrix $\mathbf{A}$ is stable just in case $\mathbf{A}^t\mathbf{y} \to \mathbf{0}$ for all $\mathbf{y} \in \mathbb{R}^n$.

Outline

Characteristic Roots and Eigenvalues

Recall the characteristic equation $|\mathbf{A} - \lambda \mathbf{I}| = 0$.

It is a polynomial equation of degree n in the unknown scalar λ.

By the fundamental theorem of algebra,
it has a set $\{\lambda_1, \lambda_2, \ldots, \lambda_n\}$ of n characteristic roots,
some of which may be repeated.

These roots may be real,
or appear in conjugate pairs $\lambda = \alpha \pm i\beta \in \mathbb{R}$ where $\alpha, \beta \in \mathbb{R}$.

Because they are roots, one can factor $|\mathbf{A} - \lambda \mathbf{I}|$ as

$$|\mathbf{A} - \lambda \mathbf{I}| = (-1)^n \prod_{i=1}^{n} (\lambda - \lambda_i)$$

When λ solves $|\mathbf{A} - \lambda \mathbf{I}| = 0$, there is a non-trivial eigenspace
of eigenvectors $\mathbf{x} \neq \mathbf{0}$ that solve the equation $\mathbf{A}\mathbf{x} = \lambda \mathbf{x}$.

Then λ is an eigenvalue.

Linearly Independent Eigenvectors

Theorem

Let $\mathbf{A}$ *be an* $n \times n$ *matrix,*
with a collection $\lambda_1, \lambda_2, \ldots, \lambda_m$ *of* $m \leq n$ *distinct eigenvalues.*

Suppose the non-zero vectors $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m$ *in* $\mathbb{R}^n$
are eigenvalues satisfying $\mathbf{A}\mathbf{u}_k = \lambda_k \mathbf{u}_k$ *for* $k = 1, 2, \ldots, m$.

Then the set $\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m\}$ *must be linearly independent.*

We prove the result by induction on m.

Note that when $m = 1$, because of the requirement that $\mathbf{u}_1 \neq \mathbf{0}$,
the set $\{\mathbf{u}_1\}$ with just one eigenvector is linearly independent.

As the induction hypothesis,
suppose that $\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_{m-1}\}$ is linearly independent.

Proof by Induction: Initial Argument

Suppose that the linear combination $\mathbf{x} = \sum_{k=1}^{m-1} \alpha_k \mathbf{u}_k$ of the linearly independent subset $\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_{m-1}\}$ of $m - 1$ vectors satisfies $\mathbf{A}\mathbf{x} = \lambda_m \mathbf{x}$.

Note that $\mathbf{A}\mathbf{x} = \sum_{k=1}^{m-1} \alpha_k \mathbf{A}\mathbf{u}_k = \sum_{k=1}^{m-1} \lambda_k \alpha_k \mathbf{u}_k$, whereas $\lambda_m \mathbf{x} = \sum_{k=1}^{m-1} \lambda_m \alpha_k \mathbf{u}_k$.

Because $\mathbf{A}\mathbf{x} = \lambda_m \mathbf{x}$, subtracting the second equation from the first gives

$$\mathbf{0} = \sum_{k=1}^{m-1} (\lambda_k - \lambda_m) \alpha_k \mathbf{u}_k$$

Then the induction hypothesis of linear independence implies that for $k = 1, \ldots, m - 1$ one has $(\lambda_k - \lambda_m)\alpha_k = 0$.

For $k = 1, \ldots, m - 1$, because $\lambda_k \neq \lambda_m$, one $\alpha_k = 0$.

So for any $\mathbf{x} \in \mathbb{R}^n$, we have proved that $\mathbf{x} = \sum_{k=1}^{m-1} \alpha_k \mathbf{u}_k$ and $\mathbf{A}\mathbf{x} = \lambda_m \mathbf{x}$ jointly imply that $\mathbf{x} = \mathbf{0}$.

Proof by Induction: The Contrapositive

To repeat, for any $\mathbf{x} \in \mathbb{R}^n$, we have proved that $\mathbf{x} = \sum_{k=1}^{m-1} \alpha_k \mathbf{u}_k$ and $\mathbf{A}\mathbf{x} = \lambda_m \mathbf{x}$ jointly imply that $\mathbf{x} = \mathbf{0}$.

The contrapositive is that $\mathbf{A}\mathbf{x} = \lambda_m \mathbf{x}$ and $\mathbf{x} \neq \mathbf{0}$ jointly imply that $\mathbf{x} \neq \sum_{k=1}^{m-1} \alpha_k \mathbf{u}_k$ for any list of scalars $(\alpha_k)_{k=1}^{m-1}$.

Hence $\mathbf{A}\mathbf{x} = \lambda_m \mathbf{x}$ and $\mathbf{x} \neq \mathbf{0}$ jointly imply that $\mathbf{x}$ must be linearly independent of $\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_{m-1}\}$.

Because $\mathbf{A}\mathbf{u}_m = \lambda_m \mathbf{u}_m$ and $\mathbf{u}_m \neq \mathbf{0}$, it follows that $\mathbf{u}_m$ is linearly independent of $\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_{m-1}\}$.

This completes the proof by induction on m. $\square$

An Eigenvector Matrix

Suppose the equation $|\mathbf{A} - \lambda \mathbf{I}| = 0$ has n distinct roots.

We remark that this holds for the generic $n \times n$ matrix $\mathbf{A}$.

In this case there are n distinct eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$.

Define the $n \times n$ eigenvector matrix $\mathbf{V} = (\mathbf{u}_j)_{j=1}^n$
whose columns are the matching set of non-zero eigenvectors.

By definition of eigenvalue and eigenvector,
for $j = 1, 2, \ldots, n$ one has $\mathbf{A}\mathbf{u}_j = \lambda_j \mathbf{u}_j$.

The j column of the $n \times n$ matrix $\mathbf{AV}$ is $\mathbf{A}\mathbf{u}_j$, which equals $\lambda_j \mathbf{u}_j$.

But with $\mathbf{\Lambda} := \mathbf{diag}(\lambda_1, \lambda_2, \ldots, \lambda_n)$, the elements of $\mathbf{V\Lambda}$ satisfy

$$(\mathbf{V\Lambda})_{ij} = \sum\nolimits_{k=1}^{n} (\mathbf{V})_{ik} \delta_{kj} \lambda_j = (\mathbf{V})_{ij} \lambda_j = \lambda_j (\mathbf{u}_j)_i$$

It follows that $\mathbf{AV} = \mathbf{V\Lambda}$ because the elements are all equal.

Diagonalization

Recall the hypothesis that the $n \times n$ matrix $\mathbf{A}$
has a full set of n distinct eigenvalues.

We have just proved this hypothesis implies that the list $(\mathbf{u}_j)_{j=1}^n$ of n associated eigenvectors must form a linearly independent set.

Hence the eigenvector matrix $\mathbf{V}$ is invertible.

We proved on the last slide that $\mathbf{AV} = \mathbf{V\Lambda}$.

Pre-multiplying this equation by $\mathbf{V}^{-1}$ yields $\mathbf{V}^{-1}\mathbf{AV} = \mathbf{\Lambda}$.

This expression is called a diagonalization of $\mathbf{A}$.

Furthermore, post-multiplying $\mathbf{AV} = \mathbf{V\Lambda}$ by the inverse matrix $\mathbf{V}^{-1}$ yields $\mathbf{A} = \mathbf{V\Lambda V}^{-1}$.

This is a decomposition of $\mathbf{A}$ into the product of:

1. the eigenvector matrix $\mathbf{V}$;
2. the diagonal eigenvalue matrix $\mathbf{\Lambda}$;
3. the inverse eigenvector matrix $\mathbf{V}^{-1}$.

Uncoupling via Diagonalization

Consider the matrix difference equation $\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1}$
for $t = 1, 2, \ldots$, with $\mathbf{x}_0$ given.

Consider the case when the $n \times n$ matrix $\mathbf{A}$ has distinct eigenvalues.

We use the invertibility of the eigenvector matrix
to define a new vector $\mathbf{y}_t = \mathbf{V}^{-1}\mathbf{x}_t$ for each t.

This new vector satisfies the transformed matrix difference equation

$$\mathbf{y}_t = \mathbf{V}^{-1}\mathbf{x}_t = \mathbf{V}^{-1}\mathbf{A}\mathbf{V}\mathbf{y}_{t-1}$$

The diagonalization $\mathbf{V}^{-1}\mathbf{A}\mathbf{V} = \mathbf{\Lambda}$ reduces this equation
to the uncoupled matrix difference equation $\mathbf{y}_t = \mathbf{\Lambda}\mathbf{y}_{t-1}$
with initial condition $\mathbf{y}_0 = \mathbf{V}^{-1}\mathbf{x}_0$.

Its solution is obviously $\mathbf{y}_t = \mathbf{\Lambda}^t\mathbf{V}^{-1}\mathbf{x}_0$
and so $\mathbf{x}_t = \mathbf{V}\mathbf{y}_t = \mathbf{V}\mathbf{\Lambda}^t\mathbf{V}^{-1}\mathbf{x}_0$.

Note that $\mathbf{\Lambda}^t = [\mathbf{diag}(\lambda_1, \ldots, \lambda_n)]^t = \mathbf{diag}(\lambda_1^t, \ldots, \lambda_n^t)$.