Lecture Notes 7: Dynamic Equations
Part C: Linear Difference Equation Systems

Peter J. Hammond

latest revision 2024 September 16th
typeset from `dynEqLects24C.tex`

Lecture Outline

Systems of Linear Difference Equations

Complementary, Particular, and General Solutions

Constant Coefficient Matrix

Some Particular Solutions

Diagonalizing a Non-Symmetric Matrix

Uncoupling via Diagonalization

Stability of Linear Systems

Stability of Non-Linear Systems

Systems of Linear Difference Equations

Many empirical economic models involve
simultaneous time series for several different variables.

Consider a first-order linear difference equation

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

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 $n \times n$ 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 $\mathbf{f}_t$ $(t = 0, 1, 2, \ldots)$,
the unique solution can be expressed as

$$\mathbf{x}_t = \mathbf{P}_{t,0} \mathbf{x}_0 + \sum_{k=1}^{t} \mathbf{P}_{t,k} \mathbf{f}_{k-1}$$

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 $\mathbf{x}_1 = \mathbf{A}_0\mathbf{x}_0 + \mathbf{f}_0$ 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

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

This formula matches the formula

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

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

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

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

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

Equating the matrix coefficients of $\mathbf{x}_0$ and of each $\mathbf{f}_{k-1}$ in this equation implies that for general t one has

$$\mathbf{P}_{t+1,k} = \mathbf{A}_t\mathbf{P}_{t,k} \text{ for } k = 0, 1, \ldots, t, \text{ with } \mathbf{P}_{t+1,t+1} = \mathbf{I}$$

Matrix Solution, II

The equation $\mathbf{P}_{t+1,k} = \mathbf{A}_t \mathbf{P}_{t,k}$ for $k = 0, 1, \ldots, t$ implies that

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

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

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

Inserting these into our formula

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

implies that

$$\mathbf{x}_t = \left(\prod\nolimits_{s=1}^{t} \mathbf{A}_{t-s}\right) \mathbf{x}_0 + \sum\nolimits_{k=1}^{t} \left(\prod\nolimits_{s=1}^{t-k} \mathbf{A}_{t-s}\right) \mathbf{f}_{k-1}$$

Lecture Outline

Complementary Solutions to the Homogeneous Equation

We are considering the general first-order linear difference equation

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

in $\mathbb{R}^n$, where each $\mathbf{A}_t$ is an $n \times n$ matrix.

The associated homogeneous equation takes the form

$$\mathbf{x}_t - \mathbf{A}_t \mathbf{x}_{t-1} = \mathbf{0} \quad (\text{for all } t \in \mathbb{N})$$

Its general solution is the n-dimensional linear subspace of functions $\mathbb{N} \ni t \mapsto \mathbf{x}_t \in \mathbb{R}^n$ satisfying

$$\mathbf{x}_t = \left(\prod_{s=1}^{t} \mathbf{A}_s \right) \mathbf{x}_0 \quad (\text{for all } t \in \mathbb{N})$$

where $\mathbf{x}_0 \in \mathbb{R}^n$ is an arbitrary constant vector.

From Particular to General Solutions

The homogeneous equation takes the form

$$\mathbf{x}_{t+1} - \mathbf{A}_t \mathbf{x}_t = \mathbf{0}$$

An associated inhomogeneous equation takes the form

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

for a general vector forcing term $\mathbf{f}_t \in \mathbb{R}^n$.

Let $\mathbf{x}_t^P$ denote a particular solution of the inhomogeneous equation and $\mathbf{x}_t^G$ any alternative general solution of the same equation.

Our assumptions imply that, for each $t = 1, 2, \ldots$, one has

$$\mathbf{x}_{t+1}^P - \mathbf{A}_t \mathbf{x}_t^P = \mathbf{f}_t \quad \text{and} \quad \mathbf{x}_{t+1}^G - \mathbf{A}_t \mathbf{x}_t^G = \mathbf{f}_t$$

Subtracting the first equation from the second implies that

$$\mathbf{x}_{t+1}^G - \mathbf{x}_{t+1}^P - \mathbf{A}_t (\mathbf{x}_t^G - \mathbf{x}_t^P) = \mathbf{0}$$

This shows that $\mathbf{x}_t^H := \mathbf{x}_t^G - \mathbf{x}_t^P$ solves the homogeneous equation.

Characterizing the General Solution

So the general solution $\mathbf{x}_t^G$
of the inhomogeneous equation $\mathbf{x}_{t+1} - \mathbf{A}_t \mathbf{x}_t = \mathbf{f}_t$
with forcing term $\mathbf{f}_t$ is the sum $\mathbf{x}_t^P + \mathbf{x}_t^H$ of

- any particular solution $\mathbf{x}_t^P$ of the inhomogeneous equation;
- the general complementary solution $\mathbf{x}_t^H$
  of the homogeneous equation $\mathbf{x}_{t+1} - \mathbf{A}_t \mathbf{x}_t = \mathbf{0}$.

Linearity in the Forcing Term

Theorem

Suppose that $\mathbf{x}_t^P$ and $\mathbf{y}_t^P$ are particular solutions of the two respective difference equations

$$\mathbf{x}_{t+1} - \mathbf{A}_t \mathbf{x}_{t-1} = \mathbf{d}_t \quad \textit{and} \quad \mathbf{y}_{t+1} - \mathbf{A}_t \mathbf{y}_{t-1} = \mathbf{e}_t$$

Then, for any scalars α and β,
the linear combination $\mathbf{z}_t^P := \alpha \mathbf{x}_t^P + \beta \mathbf{y}_t^P$ is a particular solution of the equation $\mathbf{z}_{t+1} - \mathbf{A}_t \mathbf{z}_{t-1} = \alpha \mathbf{d}_t + \beta \mathbf{e}_t$.

This can be proved by routine algebra. □

Consider any equation of the form $\mathbf{x}_{t+1} - \mathbf{A}_t \mathbf{x}_{t-1} = \mathbf{f}_t$ whose right-hand side is a linear combination $\mathbf{f}_t = \sum_{k=1}^n \alpha_k \mathbf{f}_t^k$ of the n forcing vectors $(\mathbf{f}_t^1, \ldots, \mathbf{f}_n^1)$.

The theorem implies that a particular solution is the corresponding linear combination $\mathbf{x}_t^P = \sum_{k=1}^n \alpha_k \mathbf{x}_t^{Pk}$ of particular solutions to the n equations $\mathbf{x}_{t+1} - \mathbf{A}_t \mathbf{x}_{t-1} = \mathbf{f}_t^k$.

Lecture Outline

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 $\mathbf{F}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{f}$.

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

In our previous formula, products like $\prod_{s=1}^{t-k} \mathbf{A}_{t-s}$
reduce to powers $\mathbf{A}^{t-k}$.

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

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

$$\mathbf{x}_t = \mathbf{A}^t \mathbf{x}_0 + \sum\nolimits_{k=1}^{t} \mathbf{A}^{t-k} \mathbf{f}$$

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 trick yields the equation $(1 - a)s_t = 1 - a^t$.

Hence $s_t = (1 - a)^{-1}(1 - a^t)$ 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 the two matrix equations $(\mathbf{I} - \mathbf{A})\mathbf{S}_t = \mathbf{I} - \mathbf{A}^t = \mathbf{S}_t(\mathbf{I} - \mathbf{A})$.

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

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

So the previous solution $\mathbf{x}_t = \mathbf{A}^t\mathbf{x}_0 + \sum_{k=1}^{t} \mathbf{A}^{t-k}\mathbf{f}$ reduces to

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

Lecture Outline

First-Order Linear Equation with a Constant Matrix

Recall that the solution
to the general first-order linear equation $\mathbf{x}_t - \mathbf{A}_t x_{t-1} = \mathbf{f}_t$
takes the form

$$\mathbf{x}_t = \left(\prod_{s=1}^{t} \mathbf{A}_{t-s}\right) \mathbf{x}_0 + \sum_{k=1}^{t} \left(\prod_{s=1}^{t-k} \mathbf{A}_{t-s}\right) \mathbf{f}_{k-1}$$

From now on, we restrict attention
to a constant coefficient matrix $\mathbf{A}_t = \mathbf{A}$, independent of t.

Then the solution reduces to

$$\mathbf{x}_t = \mathbf{A}^t \mathbf{x}_0 + \sum_{s=1}^{t} \mathbf{A}^{t-s} \mathbf{f}_{s-1}$$

Indeed, this is easily verified by induction.

One particular solution, of course, comes from taking $\mathbf{x}_0 = \mathbf{0}$,
implying that

$$\mathbf{x}_t = \sum_{s=1}^{t} \mathbf{A}^{t-s} \mathbf{f}_{s-1}$$

Now we will start to analyse this particular solution
for some special forcing terms $\mathbf{f}_t$.

Special Case

The special case we consider is when there exists a fixed vector $\mathbf{f}_0 \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ such that $\mathbf{x}_t - \mathbf{A}_t x_{t-1} = \mathbf{f}_t = \mu^t \mathbf{f}_0$ for the discrete exponential or power function $\mathbb{Z}_+ \ni t \mapsto \mu^t \in \mathbb{R}$.

Then the particular solution satisfying $\mathbf{x}_0 = \mathbf{0}$ is $\mathbf{x}_t = \mathbf{S}_t \mathbf{f}_0$ where $\mathbf{S}_t := \sum_{k=1}^{t} \mu^{k-1} \mathbf{A}^{t-k}$.

Note that

$$\mathbf{S}_t(\mathbf{A} - \mu \mathbf{I}) = \sum_{k=1}^{t} (\mu^{k-1} \mathbf{A}^{t-k+1} - \mu^k \mathbf{A}^{t-k}) = \mathbf{A}^t - \mu^t \mathbf{I}$$

We ignore the degenerate case when μ is an eigenvalue of $\mathbf{A}$.

Otherwise, when μ is not an eigenvalue of $\mathbf{A}$, so $\mathbf{A} - \mu \mathbf{I}$ is non-singular, it follows that

$$\mathbf{S}_t = (\mathbf{A}^t - \mu^t \mathbf{I})(\mathbf{A} - \mu \mathbf{I})^{-1}$$

Then the particular solution we are looking for takes the form

$$\mathbf{x}_t^P = (\mathbf{A}^t - \mu^t \mathbf{I})\mathbf{f}^*$$

for the particular fixed vector $\mathbf{f}^* := (\mathbf{A} - \mu \mathbf{I})^{-1} \mathbf{f}_0$.

Lecture 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 $\lambda = \alpha \pm i\beta \in \mathbb{C}$ where $\alpha, \beta \in \mathbb{R}$.

Because the λ_i are characteristic roots, one has

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

When λ solves $|\mathbf{A} - \lambda \mathbf{I}| = 0$, there is a non-trivial eigenspace E_λ
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

In the matrix algebra lectures, we proved this result:

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 $\mathbb{C}^n$ are corresponding eigenvectors satisfying

$$\mathbf{A}\mathbf{u}_k = \lambda_k \mathbf{u}_k \text{ for } k = 1, 2, \ldots, m$$

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

We also discussed similar and diagonalizable matrices.

An Eigenvector Matrix

Suppose the $n \times n$ matrix $\mathbf{A}$ has the maximum possible number of n linearly independent eigenvectors, namely $\{\mathbf{u}_j\}_{j=1}^n$.

A sufficient, but not necessary, condition for this is that $|\mathbf{A} - \lambda \mathbf{I}| = 0$ has n distinct characteristic roots.

Define the $n \times n$ eigenvector matrix $\mathbf{V} = (\mathbf{u}_j)_{j=1}^n$ whose columns are the linearly independent 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 $\boldsymbol{\Lambda} := \mathbf{diag}(\lambda_1, \lambda_2, \ldots, \lambda_n)$, the elements of $\boldsymbol{\Lambda}$ satisfy $(\boldsymbol{\Lambda})_{kj} = \delta_{kj}\lambda_j$.

So the elements of $\mathbf{V}\boldsymbol{\Lambda}$ satisfy

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

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

Diagonalization

Recall the hypothesis that the $n \times n$ matrix $\mathbf{A}$
has n linearly independent eigenvectors$\{\mathbf{u}_j\}_{j=1}^n$.

So 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 $\mathbf{V}^{-1}\mathbf{AV} = \mathbf{\Lambda}$,
which gives 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 $\mathbf{V}^{-1}$.

A Non-Diagonalizable 2×2 Matrix

Example

The non-symmetric matrix $\mathbf{A} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ cannot be diagonalized.

Its characteristic equation is $0 = |\mathbf{A} - \lambda \mathbf{I}| = \begin{vmatrix} -\lambda & 1 \\ 0 & -\lambda \end{vmatrix} = \lambda^2$.

It follows that $\lambda = 0$ is the unique eigenvalue.

The eigenvalue equation is $0 \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_2 \\ 0 \end{pmatrix}$
or $x_2 = 0$, whose only solutions take the form $x_2 \, (1, 0)^\top$.

Thus, every eigenvector is a non-zero multiple
of the column vector $(1, 0)^\top$.

This makes it impossible to find
any set of two linearly independent eigenvectors.

A Non-Diagonalizable $n \times n$ Matrix: Specification

The following $n \times n$ matrix also has a unique eigenvalue, whose eigenspace is of dimension 1.

Example

Consider the non-symmetric $n \times n$ matrix $\mathbf{A}$ whose elements in the first $n-1$ rows satisfy $a_{ij} = \delta_{i,j-1}$ for $i = 1, 2, \ldots, n-1$ but whose last row is $\mathbf{0}^\top$.

Such a matrix is upper triangular, and takes the special form

$$\mathbf{A} = \begin{pmatrix} 0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & 1 \\ 0 & 0 & 0 & \ldots & 0 \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{n-1} & \mathbf{I}_{n-1} \\ 0 & \mathbf{0}_{n-1}^\top \end{pmatrix}$$

in which the elements in the first $n-1$ rows and last $n-1$ columns make up the identity matrix.

A Non-Diagonalizable $n \times n$ Matrix: Analysis

Because $\mathbf{A} - \lambda \mathbf{I}$ is also upper triangular,
its characteristic equation is $0 = |\mathbf{A} - \lambda \mathbf{I}| = (-\lambda)^n$.

This has $\lambda = 0$ as an n-fold repeated root.

So $\lambda = 0$ is the unique eigenvalue.

The eigenvalue equation $\mathbf{A}\mathbf{x} = \lambda \mathbf{x}$ with $\lambda = 0$
takes the form $\mathbf{A}\mathbf{x} = \mathbf{0}$ or

$$0 = \sum_{j=1}^{n} \delta_{i,j-1} x_j = x_{i+1} \quad (i = 1, 2, \ldots, n-1)$$

with an extra nth equation of the form $0 = 0$.

The only solutions take the form $x_j = 0$ for $j = 2, \ldots, n$,
with x_1 arbitrary.

So all the eigenvectors of $\mathbf{A}$ are non-zero multiples
of the first canonical basis vector $\mathbf{e}_1 = (1, 0, \ldots, 0)^\top$.

This implies that there is just one eigenspace, of dimension 1.

Lecture Outline

Systems of Linear Difference Equations

Complementary, Particular, and General Solutions

Constant Coefficient Matrix

Some Particular Solutions

Diagonalizing a Non-Symmetric Matrix

Uncoupling via Diagonalization

Stability of Linear Systems

Stability of Non-Linear Systems

Uncoupling via Diagonalization

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

The extra forcing term $\mathbf{f}_t$ makes the equation inhomogeneous (unless $\mathbf{f}_t = \mathbf{0}$ for all t).

Consider the case when the $n \times n$ matrix $\mathbf{A}$
has n distinct eigenvalues,
or at least a set of n linearly independent eigenvectors
making up the columns of an invertible eigenvector matrix $\mathbf{V}$.

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

$$\mathbf{y}_t = \mathbf{V}^{-1}\mathbf{x}_t = \mathbf{V}^{-1}(\mathbf{A}\mathbf{x}_{t-1} + \mathbf{f}_t) = \mathbf{V}^{-1}\mathbf{A}\mathbf{V}\mathbf{y}_{t-1} + \mathbf{e}_t$$

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

Transforming the Uncoupled Equations

Consider the uncoupled matrix difference equation $\mathbf{y}_t = \mathbf{\Lambda} \mathbf{y}_{t-1} + \mathbf{e}_t$ where $\mathbf{\Lambda} = \mathbf{diag}(\lambda_1, \ldots, \lambda_n)$.

Note that, if there is any i for which $\lambda_i = 0$,
then the solution $\mathbf{y}_t = (y_{ti})_{i=1}^n$
must satisfy $y_{ti} = e_{ti}$ for all $t = 1, 2, \ldots$.

So we eliminate all i such that $\lambda_i = 0$,
and assume from now on that $\lambda_i \neq 0$ for all i.

This assumption ensures that $\mathbf{\Lambda}^{-1}$ exists.

This allows us to define the transformed vector $\mathbf{z}_t := \mathbf{\Lambda}^{-t} \mathbf{y}_t$ where

$$\mathbf{\Lambda}^{-t} = [\mathbf{diag}(\lambda_1, \ldots, \lambda_n)]^{-t} = \mathbf{diag}(\lambda_1^{-t}, \ldots, \lambda_n^{-t}) = (\mathbf{\Lambda}^{-1})^t$$

With this transformation, evidently

$$\mathbf{z}_t = \mathbf{\Lambda}^{-t} \mathbf{y}_t = \mathbf{\Lambda}^{-t}(\mathbf{\Lambda} \mathbf{y}_{t-1} + \mathbf{e}_t) = \mathbf{\Lambda}^{1-t} \mathbf{y}_{t-1} + \mathbf{\Lambda}^{-t} \mathbf{e}_t = \mathbf{z}_{t-1} + \mathbf{w}_t$$

where $\mathbf{w}_t$ is the transformed forcing term $\mathbf{\Lambda}^{-t} \mathbf{e}_t$.

The Decoupled Solution

The solution of $\mathbf{z}_t = \mathbf{z}_{t-1} + \mathbf{w}_t$ is obviously

$$\mathbf{z}_t = \mathbf{z}_0 + \sum_{s=1}^{t} \mathbf{w}_s$$

Inverting the previous transformation $\mathbf{z}_t = \mathbf{\Lambda}^{-t}\mathbf{y}_t$, we see that

$$\mathbf{y}_t = \mathbf{\Lambda}^t \mathbf{z}_t = \mathbf{\Lambda}^t \left(\mathbf{z}_0 + \sum_{s=1}^{t} \mathbf{w}_s \right)$$

But $\mathbf{z}_0 = \mathbf{y}_0$ and $\mathbf{w}_s = \mathbf{\Lambda}^{-s}\mathbf{e}_s$, so one has

$$\mathbf{y}_t = \mathbf{\Lambda}^t \mathbf{y}_0 + \sum_{s=1}^{t} \mathbf{\Lambda}^{t-s}\mathbf{e}_s$$

Now, each power $\mathbf{\Lambda}^k$ is the diagonal matrix $\mathbf{diag}(\lambda_1^k, \ldots, \lambda_n^k)$.

So, for each separate component y_{ti} of $\mathbf{y}_t$
and corresponding component w_{si} of $\mathbf{w}_s$,
this solution can be written in the obviously uncoupled form

$$y_{ti} = (\lambda_i)^t y_{0i} + \sum_{s=1}^{t} (\lambda_i)^{t-s} w_{si} \quad (\text{for } i = 1, 2, \ldots n)$$

The Recoupled Solution

Finally, inverting also the previous transformation $\mathbf{y}_t = \mathbf{V}^{-1}\mathbf{x}_t$, while noting that $\mathbf{e}_s = \mathbf{V}^{-1}\mathbf{f}_s$, one has

$$\mathbf{x}_t = \mathbf{V}\mathbf{y}_t = \mathbf{V}\mathbf{\Lambda}^t\mathbf{V}^{-1}\mathbf{x}_0 + \sum_{s=1}^{t} \mathbf{V}\mathbf{\Lambda}^{t-s}\mathbf{V}^{-1}\mathbf{f}_s$$

as the solution of the original equation $\mathbf{x}_t = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}\mathbf{x}_{t-1} + \mathbf{f}_t$.

Lecture Outline

Systems of Linear Difference Equations

Complementary, Particular, and General Solutions

Constant Coefficient Matrix

Some Particular Solutions

Diagonalizing a Non-Symmetric Matrix

Uncoupling via Diagonalization

Stability of Linear Systems

Stability of Non-Linear Systems

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 $\mathbb{R}^n \ni \mathbf{x} \mapsto \mathbf{F}(\mathbf{x}) \in \mathbb{R}^n$.

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 of $\Phi_1 \equiv \mathbf{F}$.

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, for $i = 1, 2, \ldots, n$,
the modulus $|\lambda_i|$ of each diagonal element λ_i satisfies $|\lambda_i| < 1$.

Recall that when $\lambda = \alpha \pm i\beta$, the modulus is $|\lambda| := \sqrt{\alpha^2 + \beta^2}$.

First 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 $\mathbf{y}_0 = \mathbf{z} = (z_1, z_2)$ 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 the solution $\mathbf{y}_t = (2^{-t} z_1, 2^t z_2)$ exhibits saddle point stability because

- starting with $z_2 = 0$ allows convergence;
- 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 $\mathbf{y} \in \mathbb{R}^n$.

Second Example: The Fibonacci Equation

Consider the Fibonacci equation $x_{t+1} = x_t + x_{t-1}$.

This has a general solution of the form $x_t = A\lambda_1^t + B\lambda_2^t$
for arbitrary constants $A, B \in \mathbb{R}$,
where $\lambda_1 = \frac{1}{2}(1 + \sqrt{5})$ and $\lambda_2 = \frac{1}{2}(1 - \sqrt{5})$ are the two roots of the quadratic characteristic equation $\lambda^2 = \lambda + 1$.

Because $|\lambda_1| > 1$ and $|\lambda_2| < 1$,
this general solution also exhibits saddle point stability.

A Condition for Stability

The solution $\mathbf{y}_t = \mathbf{A}^t \mathbf{y}_0$ of the homogeneous equation $\mathbf{y}_{t+1} = \mathbf{A}\mathbf{y}_t$ is globally stable just in case $\mathbf{A}^t \mathbf{y}_0 \to \mathbf{0}$ or $\|\mathbf{A}^t \mathbf{y}_0\| \to 0$ as $t \to \infty$, regardless of $\mathbf{y}_0$.

This holds if and only if $\mathbf{A}^t \to \mathbf{0}_{n \times n}$ in the sense that all n^2 elements of the $n \times n$ matrix $\mathbf{A}^t$ converge to 0 as $n \to \infty$.

In case the matrix $\mathbf{A}$ is the diagonal matrix $\mathbf{\Lambda} = \mathbf{diag}(\lambda_1, \lambda_2, \ldots, \lambda_n)$, stability holds if and only if $|\lambda_i| < 1$ for $i = 1, 2, \ldots, n$.

Suppose the matrix $\mathbf{A}$ is the diagonalizable matrix $\mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$, where $\mathbf{V}$ is a matrix of linearly independent eigenvectors, and the diagonal elements of the diagonal matrix $\mathbf{\Lambda}$ are eigenvalues.

Then $\mathbf{A}^t = \mathbf{V}\mathbf{\Lambda}^t\mathbf{V}^{-1} \to 0$ if and only if $\mathbf{\Lambda}^t \to 0$, which is true if and only if $|\lambda_i| < 1$ for $i = 1, 2, \ldots, n$.

The Classic Stability Condition

Definition

The $n \times n$ matrix $\mathbf{A}$ is stable just in case, as $t \to \infty$, so

1. $\mathbf{A}^t$ converges element by element to the zero matrix $\mathbf{0}_{n \times n}$;
2. or equivalently, $\mathbf{y}_t = \mathbf{A}^t \mathbf{y}_0 \to \mathbf{0}$ for all $\mathbf{y}_0 \in \mathbb{R}^n$.

Theorem

The $n \times n$ matrix $\mathbf{A}$ is stable if and only if each of its eigenvalues λ (real or complex) has modulus $|\lambda| < 1$.

We have already proved this result in case $\mathbf{A}$ is diagonalizable.

But the same stability condition applies for a general $n \times n$ matrix $\mathbf{A}$, even one that is not diagonalizable.

For such a general matrix we will only prove necessity — "only if".

Let λ^* denote the eigenvalue λ whose modulus $|\lambda|$ is largest, and let $\mathbf{x}^* \neq \mathbf{0}$ be an associated eigenvector.

In case $|\lambda^*| \geq 1$, the solution $\mathbf{x}_t = \mathbf{A}^t \mathbf{x}^* = \lambda^{*t} \, \mathbf{x}^*$ satisfies $\|\mathbf{x}_t\| = |\lambda^*|^t \|\mathbf{x}^*\| \geq \|\mathbf{x}^*\| \neq 0$, so $\mathbf{A}$ is unstable. □

Lecture Outline

Local Stability

Consider the autonomous non-linear system $\mathbf{x}_{t+1} = \mathbf{F}(\mathbf{x}_t)$ with steady state $\mathbf{x}^*$.

Let

$$\mathbf{J}(\mathbf{x}^*) = \mathbf{F}'(\mathbf{x}^*) = \left(\frac{\partial F_i}{\partial x_j} \right)_{ij} (\mathbf{x}^*)$$

denote the $n \times n$ Jacobian matrix of partial derivatives evaluated at the steady state $\mathbf{x}^*$.

Theorem

Suppose that the elements of the Jacobian matrix $\mathbf{J}(\mathbf{x}^*)$ *are continuous in a neighbourhood of the steady state* $\mathbf{x}^*$.

Let $\bar{\lambda}$ *denote the eigenvalue of* $\mathbf{J}(\mathbf{x}^*)$ *whose modulus is largest.*

The system is locally stable about the steady state $\mathbf{x}^*$*:*

$$\textit{if } |\bar{\lambda}| < 1; \quad \textit{only if } |\bar{\lambda}| \leq 1.$$

In case $|\bar{\lambda}| = 1$, the system may or may not be locally stable.

Complete Metric Spaces

Let (X, d) denote any metric space.

Definition

A Cauchy sequence $\langle x_n \rangle_{n \in \mathbb{N}}$ in X is a sequence for which,
given any $\epsilon > 0$, there exists $N_\epsilon \in \mathbb{N}$
such that $m, n > N_\epsilon \Longrightarrow d(x_m, x_n) < \epsilon$.

Definition

A metric space (X, d) is complete
just in case all its Cauchy sequences converge.

Example

Recall that one definition of the real line $\mathbb{R}$
is as the completion of the metric space $(\mathbb{Q}, d_\mathbb{Q})$,
where $\mathbb{Q}$ is the set of rational numbers,
equipped with the metric $d_\mathbb{Q}(r, r') = |r - r'|$ for all $r, r' \in \mathbb{Q}$.

That is, $(\mathbb{R}, d_\mathbb{R})$ is the smallest complete metric space
which includes the metric space $(\mathbb{Q}, d_\mathbb{Q})$.

Global Stability: Contraction Mapping Theorem

Definition

The function $X \ni x \mapsto F(x) \in X$
is a contraction mapping on the metric space (X, d)
just in case there is a positive contraction factor $K < 1$
such that $d(F(x), F(y)) \leq K\, d(x, y)$ for all $x, y \in X$.

Theorem

Suppose that $X \ni x \mapsto F(x) \in X$ is a contraction mapping on the complete metric space (X, d).

*Then for any $x_0 \in X$
the process defined by $x_t = F(x_{t-1})$ for all $t \in \mathbb{N}$
has a unique steady state $x^* \in X$ that is globally stable.*

Iteration Yields a Cauchy Sequence

Because $F : X \to X$ is a contraction mapping
with contraction factor K, and $x_t = F(x_{t-1})$ for all $t \in \mathbb{N}$,
one has $d(x_{t+1}, x_t) = d(F(x_t), F(x_{t-1})) \leq K d(x_t, x_{t-1})$.

It follows by induction on t that $d(x_{t+1}, x_t) \leq K^t d(x_1, x_0)$.

If $n > m$, then repeated application of the triangle inequality gives

$$\begin{aligned} d(x_m, x_n) &\leq \sum_{r=1}^{n-m} d(x_{m+r-1}, x_{m+r}) \\ &\leq \sum_{r=1}^{n-m} K^{m+r-1} d(x_1, x_0) \\ &= \frac{K^m - K^n}{1-K} d(x_1, x_0) < \frac{K^m}{1-K} d(x_1, x_0) \end{aligned}$$

Hence $d(x_m, x_n) < \epsilon$ provided that $K^m \leq \epsilon(1-K)/d(x_1, x_0)$ or,
since $\ln K < 0$, if $m \geq (1/\ln K)[\ln \epsilon(1-K) - \ln d(x_1, x_0)]$.

This proves that $\langle x_t \rangle_{t \in \mathbb{N}}$ is a Cauchy sequence.

Completing the Proof

Because $(x_t)_{t \in \mathbb{N}}$ is a Cauchy sequence,
the hypothesis that (X, d) is a complete metric space
implies that there is a limit point $x^* \in X$
such that $x_t \to x^*$ as $t \to \infty$.

Then, by the triangle inequality and the contraction property,

$$\begin{aligned} d(F(x^*), x^*) &\leq d(F(x^*), x_{t+1}) + d(x_{t+1}, x^*) \\ &\leq Kd(x^*, x_t) + d(x_{t+1}, x^*) \to 0 \end{aligned}$$

as $t \to \infty$, implying that $d(F(x^*), x^*) = 0$.

Because (X, d) is a metric space, it follows that $F(x^*) = x^*$,
so the limit point $x^* \in X$ is a steady state.

On the other hand, if $\bar{x} \in X$ is any steady state,
then $d(x^*, \bar{x}) = d(F(x^*), F(\bar{x})) \leq Kd(x^*, \bar{x})$.

Hence $(1 - K)d(x^*, \bar{x}) \leq 0$ which, because $K < 1$,
implies that $d(x^*, \bar{x}) \leq 0$ and so $\bar{x} = x^*$. □