PMAI LAB

Linear Algebra: Theory, Computation, and Modern Applications

Lecture 10 • Integrated Applications and Course Review

Course lens
Understand → Model → Compute → Apply → Code → Interpret

Audience
Undergraduate BS Mathematics students

Instructor
Dr. Hira Benish

From separate techniques to one complete mathematical problem-solving workflow.

1. The Course as One Connected System

Each topic answers a different question about structure, computation, or interpretation

Lecture	Mathematical focus	Main question answered
01	Linear systems and modelling	Can the real problem be represented by simultaneous linear constraints?
02	Matrices, determinants, inverses	How are transformations and equation systems computed and classified?
03	Vectors, span, and subspaces	Which outputs can combinations of available directions produce?
04	Independence, basis, dimension, rank	What information is essential, redundant, or missing?
05	Linear transformations	How does a rule map inputs to outputs while preserving linear structure?
06	Inner products, orthogonality, projections	What is the closest attainable vector in a subspace?
07	Least squares and regression	How do we fit a model when observations are noisy or inconsistent?
08	Eigenvalues and diagonalization	Which modes determine repeated or long-term transformation behaviour?
09	SVD, PCA, and compression	How can a matrix be decomposed, approximated, and reduced?

Final learning outcomes

After completing the course, students should be able to:

- translate an applied problem into vectors, matrices, and linear equations;
- select an appropriate exact, least-squares, eigenvalue, or SVD method;
- compute results by hand for small systems and with Python for realistic data;
- verify dimensions, rank, residual orthogonality, reconstruction, and stability;
- interpret numerical outputs within the original scientific or computing context; and
- communicate assumptions, limitations, and responsible conclusions.

Course-level principle
Linear algebra is not a collection of isolated formulas. It is a language for representing relationships, identifying structure, measuring approximation, and supporting decisions.

2. An End-to-End Problem-Solving Workflow

Choose the method from the structure of the problem rather than from a memorized formula

Six-step workflow

1. **Define the question.** Identify inputs, outputs, units, assumptions, and the decision to be supported.
2. **Build the representation.** Form vectors, a coefficient/design/data matrix, and any transformation or transition rule.
3. **Audit the structure.** Check dimensions, rank, independence, scaling, missing values, and conditioning.
4. **Select the method.** Use exact solving, projection, least squares, eigen-analysis, SVD, or a combination.
5. **Verify computationally.** Test equations, residuals, reconstruction, sensitivity, and numerical stability.
6. **Interpret responsibly.** Return to units and context; state what the result supports and what it does not.

Method-selection guide

Problem structure	Preferred tool	Verification
Square, consistent, full-rank system	Elimination or stable linear solver	$\mathbf{A}\boldsymbol{x} = \boldsymbol{b}$
Inconsistent or overdetermined system	Least squares / QR / SVD	$\mathbf{A}^{\mathsf{T}}\boldsymbol{r} = \mathbf{0}$
Closest point in a subspace	Orthogonal projection	residual orthogonal to subspace
Repeated transformation or dynamics	Eigenvalues and eigenvectors	$\mathbf{A}\boldsymbol{v} = \lambda\boldsymbol{v}$
Low-rank structure or compression	Truncated SVD	energy retained and reconstruction error
Correlated multivariable data	Centred PCA through SVD	explained-variance ratio

Dimension discipline
Before multiplying matrices, write their dimensions. If $\mathbf{A} \in \mathbb{R}^{m\times n}$ and $\boldsymbol{x} \in \mathbb{R}^{n}$, then $\mathbf{A}\boldsymbol{x} \in \mathbb{R}^{m}$. Many modelling and coding errors are dimension errors.

Numerical discipline
Prefer `solve`, `lstsq`, QR, or SVD to explicitly forming an inverse. Mathematical equivalence does not guarantee equal numerical stability.

3. Integrated Case Study: Smart-City Environmental Monitoring

Predict a pollution index and reduce correlated sensor variables

Five monitoring periods record a temperature-above-baseline index x_1, a traffic index x_2, and a pollution index y:

Period	1	2	3	4	5
x_1	1	2	3	4	5
x_2	2	1	4	3	5
y	3	3	7	6	9

Question A: predictive model

Assume

$$\widehat{y} = \beta_0 + \beta_1 x_1 + \beta_2 x_2.$$

The design matrix, coefficient vector, and response vector are

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & 2 \\ 1 & 2 & 1 \\ 1 & 3 & 4 \\ 1 & 4 & 3 \\ 1 & 5 & 5 \end{bmatrix}, \qquad \boldsymbol{\beta} = \begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \end{pmatrix}, \qquad \boldsymbol{y} = \begin{pmatrix} 3 \\ 3 \\ 7 \\ 6 \\ 9 \end{pmatrix}.$$

Because there are five observations and three unknown coefficients, the system is overdetermined. The columns are independent, so rank$(\mathbf{A}) = 3$.

Question B: dimension reduction

The predictor matrix is

$$\mathbf{X} = \begin{bmatrix} 1 & 2 \\ 2 & 1 \\ 3 & 4 \\ 4 & 3 \\ 5 & 5 \end{bmatrix}.$$

Temperature and traffic move together. PCA asks whether one combined variable can represent most of their variation.

Why several course concepts are needed

- Matrix modelling organizes observations and coefficients.
- Rank confirms identifiable regression coefficients.
- Least squares fits noisy data.
- Projection explains fitted values and residuals.
- Eigenvalues and SVD reveal principal directions.
- Python supports reliable computation and visualization.

Modelling assumption
The fitted coefficients describe association within these observations. They do not prove that temperature or traffic alone causes pollution.

4. Regression, Projection, and Prediction

Solve the overdetermined model and interpret the residual geometry

The normal equations are

$$\mathbf{A}^{\top}\mathbf{A}\widehat{\boldsymbol{\beta}} = \mathbf{A}^{\top}\boldsymbol{y},$$

with

$$\mathbf{A}^{\top}\mathbf{A} = \begin{bmatrix} 5 & 15 & 15 \\ 15 & 55 & 53 \\ 15 & 53 & 55 \end{bmatrix}, \qquad \mathbf{A}^{\top}\boldsymbol{y} = \begin{pmatrix} 28 \\ 99 \\ 100 \end{pmatrix}.$$

Solving gives

$$\widehat{\boldsymbol{\beta}} = \begin{pmatrix} 0.4333 \\ 0.6111 \\ 1.1111 \end{pmatrix}, \qquad \boxed{\widehat{y} = 0.4333 + 0.6111x_1 + 1.1111x_2.}$$

Fitted values and residuals

$$\widehat{\boldsymbol{y}} \approx \begin{pmatrix} 3.2667 \\ 2.7667 \\ 6.7111 \\ 6.2111 \\ 9.0444 \end{pmatrix}, \qquad \boldsymbol{r} = \boldsymbol{y} - \widehat{\boldsymbol{y}} \approx \begin{pmatrix} -0.2667 \\ 0.2333 \\ 0.2889 \\ -0.2111 \\ -0.0444 \end{pmatrix}.$$

The residual satisfies

$$\mathbf{A}^{\top}\boldsymbol{r} = \mathbf{0},$$

so the fitted vector is the orthogonal projection of $\boldsymbol{y}$ onto Col($\mathbf{A}$).

Model evaluation

$$\text{SSE} \approx 0.2556, \qquad \text{RMSE} \approx 0.2261, \qquad R^2 \approx 0.9906.$$

Planning prediction

For a future period with $x_1 = 6$ and $x_2 = 4$,

$$\widehat{y} = 0.4333 + 0.6111(6) + 1.1111(4) \approx 8.5444.$$

Application decision
Use 8.54 as a model-based pollution-index estimate for this predictor combination. The traffic coefficient is larger in this fitted scale, but coefficient comparison depends on measurement units and predictor correlation.

Validation caution
The sample is very small. A high R^2 on the same data used for fitting does not guarantee future accuracy; additional observations and out-of-sample validation are required.

5. PCA, Eigenvalues, and SVD

Replace two correlated predictors by a dominant combined direction

The predictor means are

$$\bar{x}_1 = 3, \qquad \bar{x}_2 = 3.$$

Therefore,

$$\mathbf{X}_c = \begin{bmatrix} -2 & -1 \\ -1 & -2 \\ 0 & 1 \\ 1 & 0 \\ 2 & 2 \end{bmatrix}.$$

The sample covariance matrix is

$$\mathbf{C} = \frac{1}{4}\mathbf{X}_c^\top \mathbf{X}_c = \begin{bmatrix} 2.5 & 2.0 \\ 2.0 & 2.5 \end{bmatrix}.$$

Eigen-analysis

The eigenvalues are

$$\lambda_1 = 4.5, \qquad \lambda_2 = 0.5,$$

with principal directions

$$\boldsymbol{v}_1 = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad \boldsymbol{v}_2 = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

The explained-variance ratios are

$$\frac{4.5}{4.5+0.5} = 90\%, \qquad \frac{0.5}{4.5+0.5} = 10\%.$$

SVD connection

If

$$\mathbf{X}_c = \mathbf{U}\Sigma\mathbf{V}^\top,$$

then

$$\sigma_1 = \sqrt{4(4.5)} = \sqrt{18} \approx 4.2426, \qquad \sigma_2 = \sqrt{4(0.5)} = \sqrt{2} \approx 1.4142.$$

The first principal-component scores are

$$\mathbf{X}_c\boldsymbol{v}_1 \approx \begin{pmatrix} -2.1213 \\ -2.1213 \\ 0.7071 \\ 0.7071 \\ 2.8284 \end{pmatrix}.$$

Dimension-reduction decision
One principal component retains 90% of predictor variance. It can serve as a compact combined temperature-traffic intensity index when the lost 10% is acceptable.

Interpretation caution
PCA maximizes variance without using the pollution response. A direction with lower predictor variance may still be important for prediction or policy interpretation.

6. Linear Algebra across Modern Fields

The same mathematical structures support applications far beyond one dataset

Field	Linear-algebra tool	Application interpretation
Computer graphics	Transformations, homogeneous matrices	Rotate, scale, translate, and project 2D or 3D objects
Robotics and navigation	Linear systems, least squares, projections	Estimate position from noisy sensors and constraints
Machine learning	Matrices, gradients, least squares, SVD	Represent data, fit parameters, reduce dimension, and stabilize computation
Networks and search	Eigenvectors and transition matrices	Identify long-run importance, centrality, or flow distribution
Signal and image processing	Orthogonal bases and truncated SVD	Denoise, compress, reconstruct, and isolate dominant modes
Finance and economics	Covariance eigenvectors and PCA	Summarize correlated variables and construct factor models
Epidemiology and ecology	Transition matrices and eigenvalues	Study growth, decay, reproduction, and equilibrium behaviour
Natural language processing	Vector spaces, inner products, low rank	Compare embeddings and discover latent semantic structure

A four-part interpretation checklist

1. **Mathematical:** Is the result algebraically correct and dimensionally valid?
2. **Computational:** Is the algorithm stable, reproducible, and verified?
3. **Domain:** What do coefficients, directions, residuals, or modes mean in context?
4. **Responsible:** Which assumptions, missing variables, biases, and consequences matter?

Global perspective
The formulas are universal, but interpretation is local. Units, data collection, cultural and geographic context, uncertainty, and decision consequences determine whether a mathematically correct model is useful.

Career connection

Linear algebra supports work in data science, AI, numerical computing, operations research, quantitative finance, computational biology, engineering simulation, geospatial analysis, computer vision, and scientific research.

7. Integrated Python Laboratory

Fit the model, verify projection geometry, perform PCA, and visualize results

```
import numpy as np
import matplotlib.pyplot as plt

# Smart-city observations
x1 = np.array([1., 2., 3., 4., 5.])  # temperature index
x2 = np.array([2., 1., 4., 3., 5.])  # traffic index
y = np.array([3., 3., 7., 6., 9.])   # pollution index

# Multiple least-squares regression
A = np.column_stack([np.ones_like(x1), x1, x2])
beta, residual_sum, rank, singular_A = np.linalg.lstsq(
    A, y, rcond=None
)
y_hat = A @ beta
residuals = y - y_hat
sse = residuals @ residuals
rmse = np.sqrt(np.mean(residuals**2))
r_squared = 1 - sse / np.sum((y - y.mean())**2)
orthogonality_check = A.T @ residuals
prediction = np.array([1., 6., 4.]) @ beta

print("Design-matrix rank:", rank)
print("Regression coefficients:", np.round(beta, 4))
print("SSE, RMSE, R-squared:",
      round(float(sse), 4), round(float(rmse), 4),
      round(float(r_squared), 4))
print("A.T @ residuals:", np.round(orthogonality_check, 10))
print("Prediction for x1=6, x2=4:", round(float(prediction), 4))

# PCA through SVD
X = np.column_stack([x1, x2])
X_centered = X - X.mean(axis=0)
U, s, Vt = np.linalg.svd(X_centered, full_matrices=False)
scores = X_centered @ Vt.T
explained_ratio = s**2 / np.sum(s**2)

print("PCA singular values:", np.round(s, 4))
print("Principal directions:\n", np.round(Vt.T, 4))
print("Explained-variance ratio:", np.round(explained_ratio, 4))

# Visual checks
fig, axes = plt.subplots(1, 2, figsize=(10, 3.6))
axes[0].scatter(y, y_hat, s=65, color="#1665D8")
axes[0].plot([y.min(), y.max()], [y.min(), y.max()],
             "--", color="#E5A823")
axes[0].set_xlabel("Observed pollution index")
axes[0].set_ylabel("Predicted pollution index")
axes[0].set_title("Observed versus predicted")

axes[1].scatter(scores[:, 0], scores[:, 1],
                s=65, color="#0C8F8F")
for i, (a, b) in enumerate(scores, start=1):
    axes[1].annotate(str(i), (a, b), xytext=(4, 4),
                     textcoords="offset points")
axes[1].set_xlabel("Principal component 1")
axes[1].set_ylabel("Principal component 2")
axes[1].set_title("Centred predictor scores")
plt.tight_layout()
plt.show()
```

Selected output

```
Design-matrix rank: 3
Regression coefficients: [0.4333 0.6111 1.1111]
SSE, RMSE, R-squared: 0.2556 0.2261 0.9906
A.T @ residuals: [0. 0. 0.]
Prediction for x1=6, x2=4: 8.5444
PCA singular values: [4.2426 1.4142]
Explained-variance ratio: [0.9 0.1]
```

8. Final Concept Check and Practice

Three concept-check MCQs

1. For an overdetermined inconsistent system, the most appropriate general method is:

A. determinant expansion

B. least squares

C. row sorting

D. scalar multiplication

2. The residual of a least-squares solution is orthogonal to:

A. Col($\mathbf{A}$)

B. every vector in $\mathbb{R}^n$

C. the response only

D. the identity matrix

3. The dominant principal direction of centred data can be found from:

A. the smallest row sum

B. the leading right singular vector

C. the determinant alone

D. the zero vector

MCQ answers: 1-B, 2-A, 3-B

Practice exercises

1. Solve $2x + y = 7$ and $x - y = 2$.
2. Find the rank of $\begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$.
3. Project $\boldsymbol{b} = (3,1)^\top$ onto the line spanned by $\boldsymbol{u} = (1,1)^\top$ and find the residual.
4. Fit a line to $(0,1), (1,2), (2,3)$ and report its residuals.
5. State the long-term stability implication for eigenvalues 0.8 and 0.4 in a diagonalizable discrete system.
6. A matrix has singular values 3 and 1. Find the energy retained by its rank-one approximation.
7. For the integrated case, report the prediction at $(x_1, x_2) = (6,4)$ and the PC1 explained variance.
8. Replace the final pollution observation 9 by 10, refit in Python, and compare the coefficients and R^2.

A compact formula map

Concept	Core relation
Linear system	$\mathbf{A}\boldsymbol{x} = \boldsymbol{b}$
Projection / least squares	$\mathbf{A}^\top(\boldsymbol{b} - \mathbf{A}\widehat{\boldsymbol{x}}) = \mathbf{0}$
Eigenpair	$\mathbf{A}\boldsymbol{v} = \lambda\boldsymbol{v}$
Diagonalization	$\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1}$
SVD	$\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^\top$
PCA variance ratio	$\sigma_j^2 / \sum_i \sigma_i^2$

9. Capstone Challenge, Answers, and Completion

Capstone scenario: smart-agriculture sensor fusion

Four periods record soil-moisture index x_1, sunlight index x_2, and crop-growth index y:

$$\mathbf{X} = \begin{bmatrix} 1 & 2 \\ 2 & 1 \\ 3 & 4 \\ 4 & 3 \end{bmatrix}, \qquad \boldsymbol{y} = \begin{pmatrix} 4 \\ 4 \\ 8 \\ 8 \end{pmatrix}.$$

Fit $\widehat{y} = \beta_0 + \beta_1 x_1 + \beta_2 x_2$, predict at $(5, 4)$, and use PCA to determine the variance retained by the first predictor component.

Brief answers to practice

Exercise	Brief answer
1	$(x, y) = (3, 1)$
2	Rank $= 1$ because the second row is twice the first
3	Projection $= (2, 2)^\mathsf{T}$; residual $= (1, -1)^\mathsf{T}$
4	$\widehat{y} = 1 + x$; all residuals are zero
5	Both modes decay because $\lvert\lambda_i\rvert < 1$; the state approaches $\mathbf{0}$
6	Rank-one energy $= 3^2/(3^2 + 1^2) = 90\%$
7	Prediction ≈ 8.5444; PC1 explains 90%
8	$\widehat{\boldsymbol{\beta}} \approx (-0.0333, 0.7222, 1.2222)^\mathsf{T}$ and $R^2 \approx 0.9850$
Capstone	$\widehat{\boldsymbol{\beta}} = (1, 1, 1)^\mathsf{T}$; prediction $= 10$; PCA singular values $2.8284, 1.4142$, so PC1 retains 80%

Final takeaways

- Start with a clear question and a valid mathematical representation.
- Use rank, dimension, orthogonality, eigenvalues, and singular values to understand structure.
- Combine hand reasoning with reliable Python computation and verification.
- Treat approximation error, stability, assumptions, and interpretation as part of the solution.
- Apply linear algebra as a transferable foundation for mathematics, science, computing, and AI.

Course completion
You now have a connected toolkit for modelling, solving, approximating, reducing, and interpreting linear problems. The next step is to apply these ideas to larger datasets, optimization, differential equations, numerical analysis, machine learning, and research.

Key terms revisited: system • matrix • vector • span • basis • rank • transformation • inner product • projection • least squares • eigenvalue • diagonalization • SVD • PCA

Suggested references: Gilbert Strang, Introduction to Linear Algebra; David C. Lay, Steven R. Lay, and Judi J. McDonald, Linear Algebra and Its Applications; NumPy linear-algebra documentation.