

Lesson 31: PCA Dimensionality Reduction

Data Science with Python – BSc Course

Data Science Program

BSc Course

45 Minutes

Previously on L30: Hierarchical Clustering

What We Learned:

- Dendrograms reveal nested cluster structure at every scale
- Ward linkage minimizes within-cluster variance
- Correlation-based clustering groups co-moving assets
- HRP uses the hierarchy for portfolio allocation



We can cluster observations and features. But what if we have 100 correlated features?

The Curse of Too Many Features

Situation: You have 100 features for each stock.

What goes wrong?

- Many features are **correlated** (P/E and earnings yield, volatility and beta)
- Models become **slow** – more features = more computation
- Noise overwhelms signal – **overfitting** with sparse high-dimensional data
- Visualization impossible beyond 3D

Complication: We can't just drop features randomly – we might lose important information hidden in the correlations.

Question: Can we compress 100 features into a handful while keeping the essential patterns?

High dimensions: slow, noisy, overfitting. We need a principled way to compress.

Learning Objectives

The Problem: 100 features, many correlated. Models are slow, noisy, and overfit. Can we compress without losing the big picture?

After this lesson, you will be able to:


- Explain PCA as finding new axes of maximum variance
- Apply PCA with sklearn (StandardScaler → PCA → fit_transform)
- Interpret scree plots and explained variance ratios
- Use loadings and biplots to understand what components mean

Finance Application: Factor discovery, risk decomposition, noise reduction

The Photo Compression Idea

Analogy: 100-Megapixel Photo → JPEG

- Raw photo stores every pixel – massive, redundant
- JPEG keeps shapes and colors, drops tiny details
- Result: 10× smaller, still looks great




PCA does this for data: compress correlated features into fewer components.

PCA = JPEG for your feature matrix. Lose some detail, keep the essence.

PCA Concept: Finding New Axes

Rotating the Coordinate System

- Original axes may be correlated – PCA finds new axes aligned with spread
- PC1: maximum variance direction. PC2: perpendicular, next most.

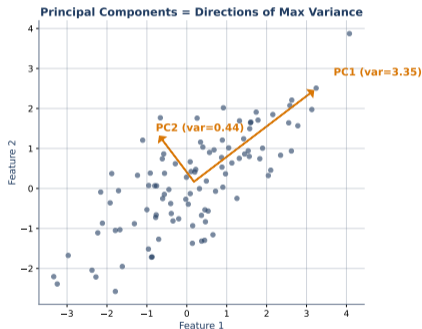


PCA rotates the coordinate system to align with the data's natural directions

Principal Component Directions

Each PC = weighted combination of original features

- $PC1 = w_1 \times \text{Return} + w_2 \times \text{Volatility} + \dots$
- Weights maximize variance; all PCs are orthogonal (uncorrelated)




Orthogonal = perpendicular = uncorrelated. Each PC captures unique information.

Projection: 2D to 1D

Dimensionality Reduction in Action

- Project each point onto PC1 – collapse 2D to 1D
- If PC1 captures 90% of variance, we keep 90% of information




Projection = shadow of data on PC1. More variance on PC1 = better shadow.

Why PCA Works

Correlation Is Compressible

- Correlated features carry overlapping information
- PCA exploits this redundancy – fewer components, same signal




High correlation = high compressibility. PCA works best on correlated features.

Eigenvectors: The Data's Natural Axes

Where Do PC Directions Come From?

- Covariance matrix encodes all feature relationships
- Its eigenvectors = PC directions; eigenvalues = variance per PC



Eigenvector = direction of stretch. Eigenvalue = amount of stretch. Larger = more important.

PCA in sklearn

Three Lines of Code

```
from sklearn.preprocessing import StandardScaler  
from sklearn.decomposition import PCA
```

1. `X_scaled = StandardScaler().fit_transform(X)`
2. `pca = PCA(n_components=3)`
3. `X_pca = pca.fit_transform(X_scaled)`

Key Attributes After Fitting:

- `pca.explained_variance_ratio_` – fraction of variance per PC
- `pca.components_` – PC directions (weights for each feature)
- `pca.n_components_` – number of components kept


Critical: Always scale first! Unscaled features with large values dominate.

StandardScaler → **PCA** → **fit_transform**. **Always scale first!**

Scree Plot: How Many Components?

Plot Eigenvalues in Decreasing Order

- Look for the “elbow” where values level off – cut there
- Named after geological scree (rubble at a mountain base)



Mountain peak = important PCs. Rubble (scree) = noise. Cut at the elbow.

Checkpoint: Variance Math

Quick Exercises – Think Before You Peek

Q1: PC1 explains 60%, PC2 explains 25%.
Together they explain ??% of variance.

Q2: Individual ratios: 50%, 25%, 15%, 7%, 3%.
How many components for 95% variance?

Answers:


- Q1: $60 + 25 = \mathbf{85\%}$ (cumulative = sum of individual ratios)
- Q2: $50+25+15+7 = 97\% \geq 95\% \Rightarrow \mathbf{4 \text{ components}}$

Cumulative variance = running sum of explained variance ratios. Simple addition!

Cumulative Explained Variance

Running Total of Information Retained

- Cumulative curve: total variance captured by first k PCs
- Read off: “How many components for 95% variance?”




Steep curve = good compression. Flat curve = data is complex, many PCs needed.

Explained Variance Ratio

Individual + Cumulative View

- Bars: each component's share. Line: cumulative sum.
- Decision rule: keep PCs until cumulative ≥ 0.95




95% variance threshold is convention. Adjust: visualization (80%), reconstruction (99%).

Loadings: What Do Components Mean?

Interpreting the “Black Box”

- Loadings = correlation between original features and each PC
- High loading = strong influence; sign shows direction




Loadings turn abstract PCs into interpretable factors. “PC1 = market exposure.”

Loadings Heatmap

Which Features Drive Which Components?

- Rows = features, columns = PCs. Dark = strong influence.
- Read column-wise to see what each PC captures.




Heatmap pattern reveals what each PC “means” in terms of original features.

Biplot: Data and Loadings Together

The Most Informative PCA Visualization

- Points: observations on PC1 vs PC2. Arrows: feature directions.
- Nearby points are similar; arrows show why.




Biplot = scores (points) + loadings (arrows). Arrow direction = feature's PC alignment.

Finance: Factor Extraction from Returns

PCA on Stock Returns = Statistical Factors

- PC1 \approx market factor (all stocks move together)
- PC2–PC3 often capture sector or style effects
- Compare: PCA is data-driven, Fama-French (L24) is theory-driven




PCA discovers latent factors from returns. Compare with Fama-French from L24.

Finance: Factor Loadings by Stock

Each Stock's Exposure to Statistical Factors

- Loadings show each stock's factor exposure
- Tech loads on PC2 (growth); utilities load oppositely (defensive)



Loadings = factor exposures. Portfolio managers use these for risk budgeting.

PCA Limitations

When PCA Struggles

1. Non-Linear Relationships

- PCA finds linear directions only
- Curved or clustered structure may be missed entirely
- Alternative: Kernel PCA, t-SNE, UMAP for nonlinear patterns

2. Scaling Sensitivity

- Unscaled data \Rightarrow PC1 dominated by largest-variance feature
- Always standardize unless features share the same units

3. Interpretability

- PCs are mathematical constructs, not natural concepts
- “PC1 = 0.35×Return + 0.28×Vol + ...” – what does it *mean*?
- Loadings and domain knowledge help, but naming is subjective


PCA is powerful but linear. Check assumptions before trusting the compression.

What's Next: The ML Assembly Line

Closing: From Components to Complete Pipelines

We now have all the building blocks:

- Preprocessing (scaling, encoding)
- Dimensionality reduction (PCA)
- Models (regression, classification, clustering)



Next lesson (L32): Putting it all together in sklearn Pipeline – reproducible, leak-free, deployable.

L32 preview: sklearn Pipeline chains Scale → PCA → Model into one object.

Hands-On Exercise (25 min)

Task: Discover Factors in Stock Returns

1. Load daily returns for 20–30 stocks (1 year of data)
2. Standardize returns with `StandardScaler`
3. Fit `PCA()` and plot a scree plot
4. How many components for 90% variance?
5. Examine PC1 loadings – what does it represent?
6. Create a biplot: project stocks to PC1 vs PC2, color by sector

Deliverable: Scree plot + biplot with sector colors.

Extension: Correlate PC1 scores with S&P 500 returns. Is PC1 the “market factor”?

Tip: Use `yfinance` to download stock prices, then `pct_change()` for returns.

Lesson Summary

Problem Solved: We can now compress many correlated features into a few uncorrelated components while preserving most of the information.

Key Takeaways:

- PCA finds orthogonal axes of maximum variance (new coordinate system)
- Scree plot + cumulative variance guide how many components to keep
- **Always standardize** before PCA – scaling matters!
- Loadings and biplots reveal what each component means
- Finance: PCA extracts statistical factors from return matrices

Next Lesson: L32 – ML Pipeline (putting it all together)

Memory: PCA = rotate to max-variance axes. Always scale. Loadings = interpretation.