AutoEncoders (DRAFT)

Mark A. Austin

University of Maryland

austin@umd.edu
ENCE 688P, Spring Semester 2026

October 17, 2026

Overview

1. Introduction to AutoEncoders
   - Basic Idea and Applications
2. Computational Foundations
   - Dimensionality Reduction
   - Principal Component Analysis
3. Applications
   - Simplest Example: PCA vs AutoEncoder
   - Linear vs Nonlinear Dimensionality Reduction
   - Handwritten Digit Recognition
4. Appendix A: Cheatsheets

AutoEncoder

Basic Idea and Applications

Autoencoders

Autoencoder neural networks use unsupervised machine learning algorithms to: (1) find compressed representations of the input data (encoder), and (2) reconstruct the original data from the compressed data (decoder).

Applications:

- Dimensionality reduction.
- Image processing (compression and denoising).
- Feature extraction; anomaly detection.
- Image generation.
- Sequence-to-sequence translation.
- Recommendation systems.

AutoEncoder Architecture

AutoEncoder (Encoder-Decoder-Reconstruction)

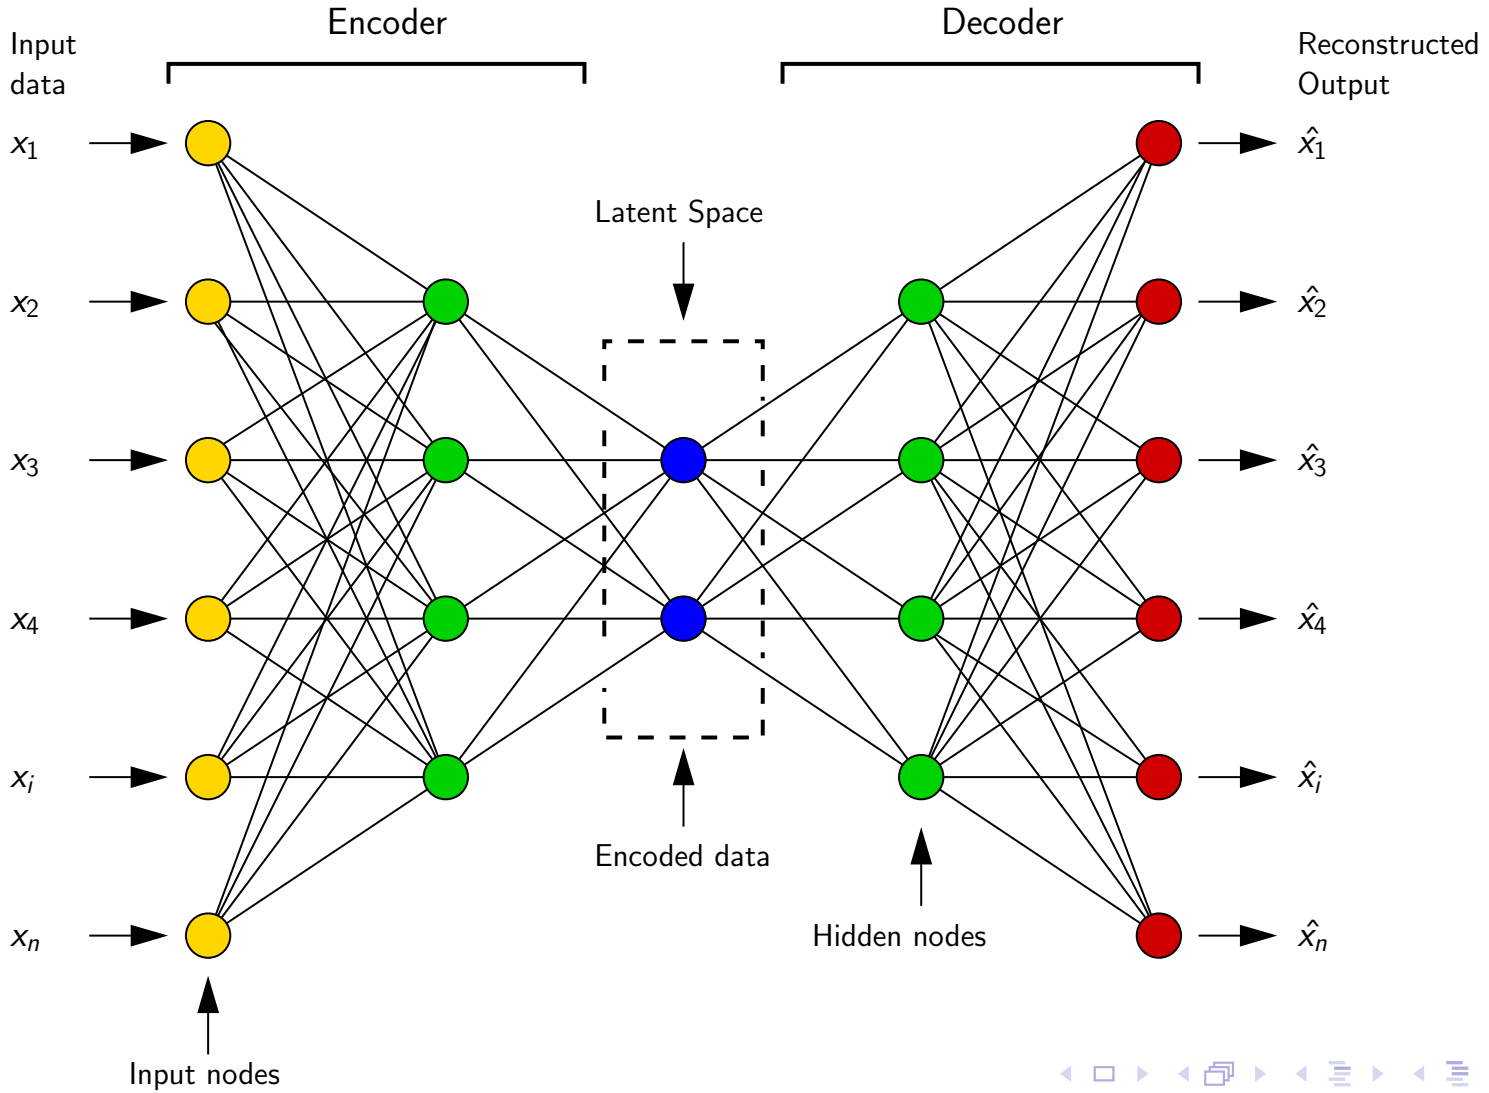

AutoEncoder Architecture

Encoder

The encoder learns how to reduce the input dimensions and compress the input data into an encoded representation.

Decoder

The decoder learns how to reconstruct the input data from the encoded representation and be as close to the input data as possible.

Latent Space

Latent space is simply a representation of compressed data in which similar points are closer together in space. This formalism is useful for learning data features and finding similar representations of data for analysis.

AutoEncoder Design and Analysis

Anomaly Detection ...

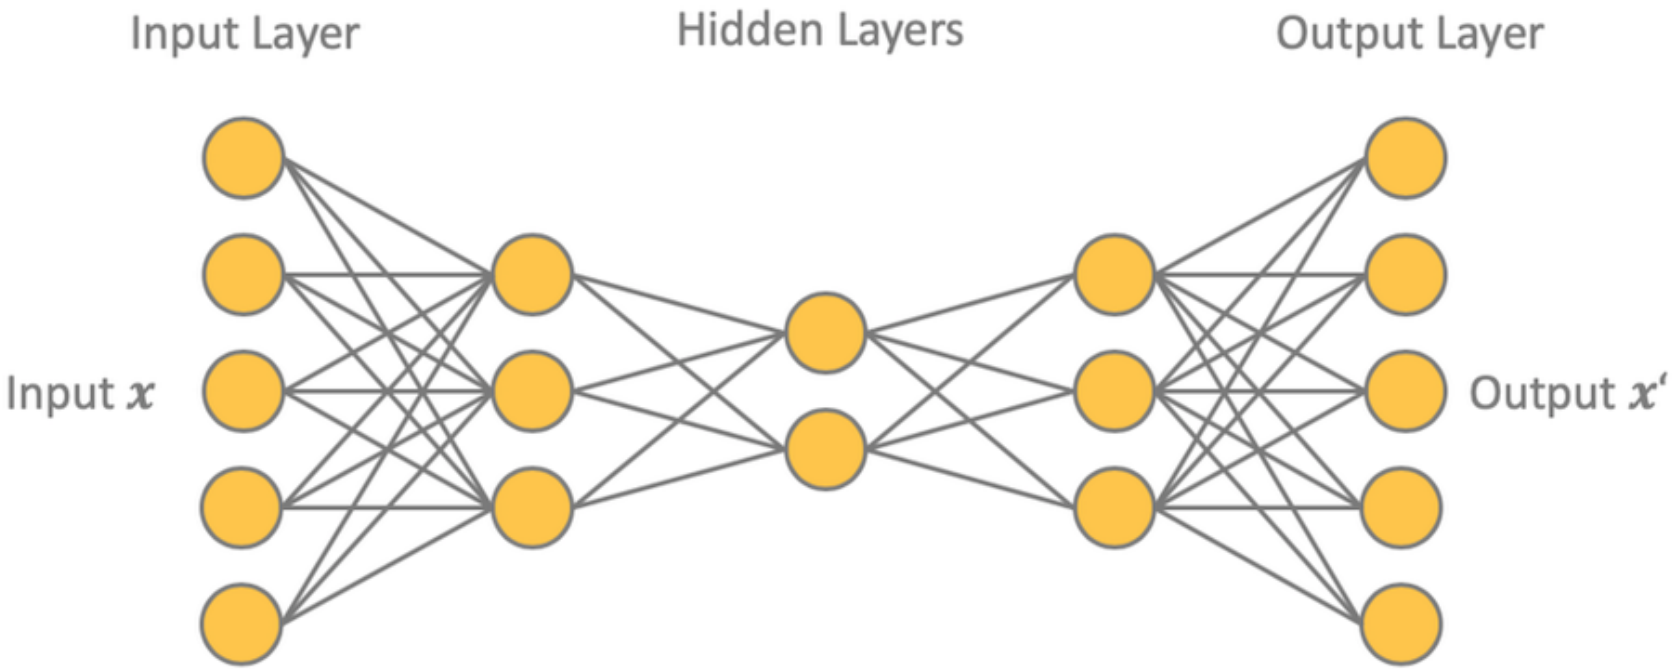

Execution of the Network:

$$\sqrt{(x_{new} - x'_{new})^2} > \delta \Rightarrow \text{anomaly}$$

AutoEncoder Design and Analysis

Ideal Requirements/Properties:

- **Tied Weights.** Weights in i-th layer of the encoder are equal to the transpose of weights in the i-th layer of the decoder, i.e., $W_i = W_i^T$.
- **Orthogonal Weights.** Weights in the encoder are orthogonal, i.e., $W_i^T.W_i = I$.
- **Uncorrelated Features.** Encoding layer outputs are not correlated.
- **Unit Norm.** The weights have unit norm, i.e.,

$$\sum_{j=1}^{p} w_{ij}^2 = 1 \text{ for } i = 1 \cdots k. \tag{1}$$

Computational Foundations

Dimensionality Reduction

Dimensionality Reduction

Strategies of dimensionality reduction involve transformation of data to new (lower) dimension in such a way that some of the dimensions can be discarded without a loss of information.

Example: Projection of Swiss Roll data in 3D to 2D ...

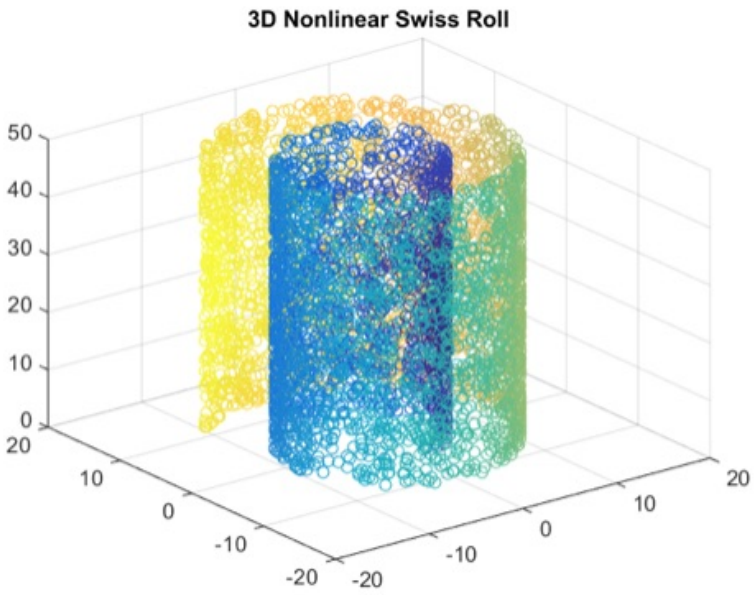

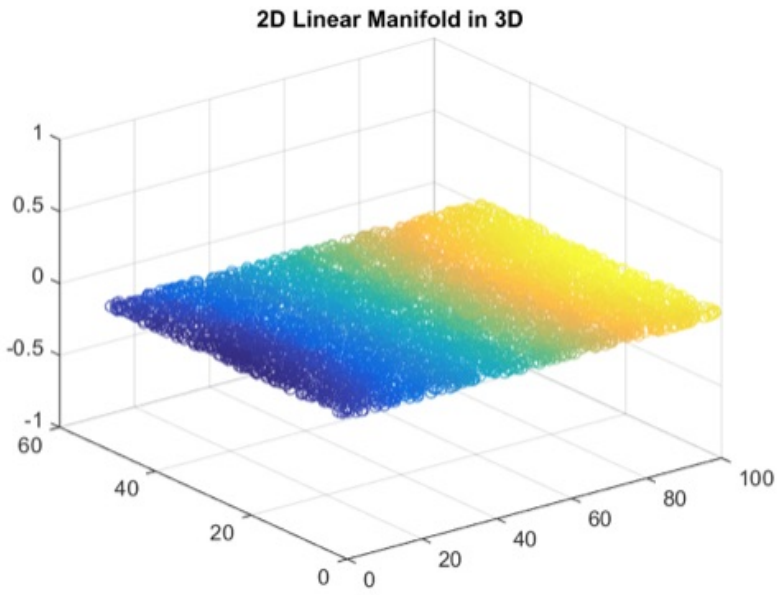

Principal Component Analysis

Principal Component Analysis

Principal component analysis is an orthogonal linear transformation that transforms mean-centered data into a new coordinate system such that the variance of the projected data is maximized along the new axes – the latter is called the first principal component.

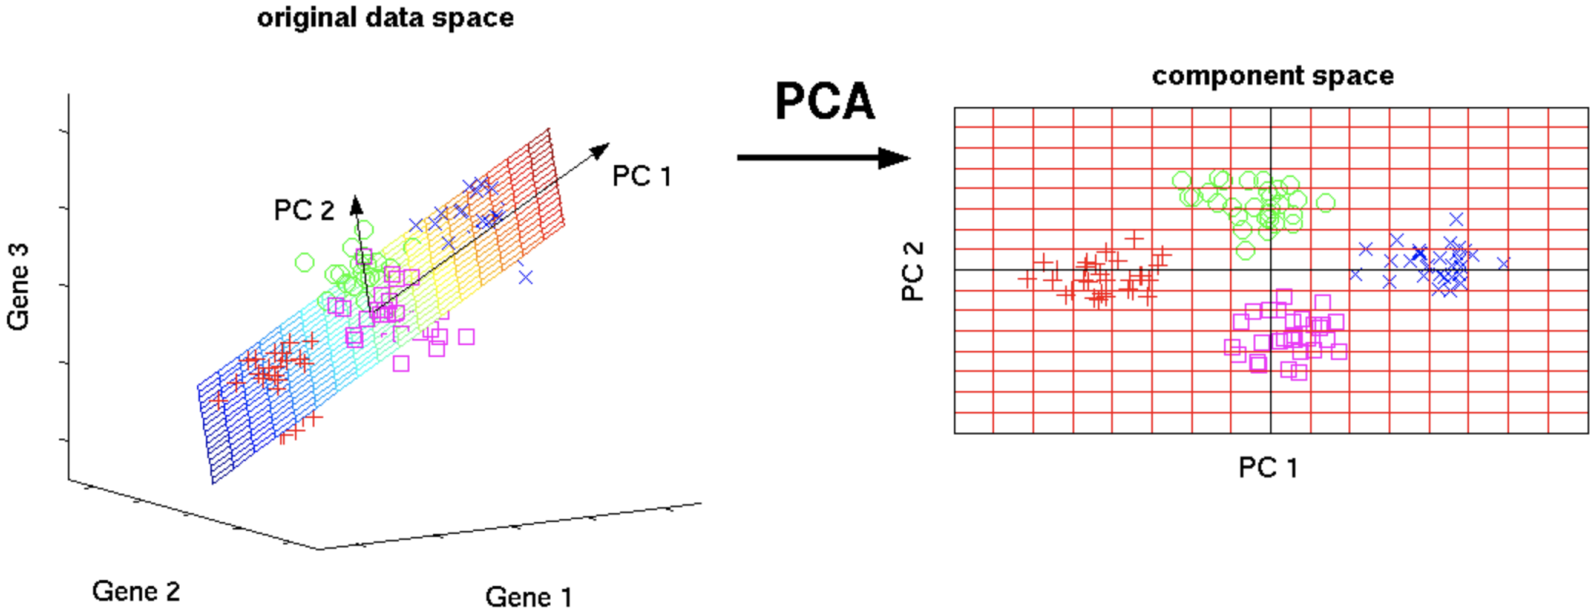

Principal Component Analysis

Interpretation of Principal Components:

- The first principal component can be defined as the direction that maximizes the variance of the projected data.
- The i-th principal component is a direction that maximizes the variance in the projected data and is orthogonal to the first i-1 principal components.

Applications:

- Exploratory data analysis.
- Dimensionality reduction.

Note: In general, dimensionality reduction loses information. PCA-based reduction procedures tend to minimize information loss.

Principal Component Analysis

Mathematical Procedure: Suppose that our dataset comprises n m-dimensional data points:

$$X = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1m} \\ x_{21} & x_{22} & \cdots & x_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \cdots & x_{nm} \end{bmatrix} \tag{2}$$

Step 1. Compute the mean value for each dimension in the dataset:

$$\bar{X} = \begin{bmatrix} \bar{x}_1 & \bar{x}_2 & \cdots & \bar{x}_m \end{bmatrix} \tag{3}$$

Principal Component Analysis

Step 2. Compute the $m \times m$ covariance matrix:

$$\text{Cov} = \begin{bmatrix} cov_{11} & cov_{12} & \cdots & cov_{1m} \\ cov_{21} & cov_{22} & \cdots & cov_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ cov_{m1} & cov_{m2} & \cdots & cov_{mm} \end{bmatrix} \tag{4}$$

where,

$$cov_{ij} = \frac{1}{n-1} \sum_{k=1}^{n} (x_{ki} - \bar{x}_i)(x_{kj} - \bar{x}_j) \tag{5}$$

Step 3. Compute the covariance matrix eigenvalues/eigenvectors:

$$[\text{Cov}]\, W = \lambda W. \tag{6}$$

Principal Component Analysis

Step 4. Sort eigenvalues by decreasing order.

Step 5. Choose k ($k \leq m$) largest eigenvalues/eigenvectors to form $m \times k$ matrix:

$$W = \begin{bmatrix} w_{11} & w_{12} & \cdots & w_{1k} \\ w_{21} & w_{22} & \cdots & w_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ w_{m1} & w_{m2} & \cdots & w_{mk} \end{bmatrix}. \tag{7}$$

Step 6. Transform raw data X onto k-dimensional subspace Y:

$$Y = X \cdot W. \tag{8}$$

Principal Component Analysis

Example 1: Define straight line segment + noisy data:

```
def NoisyLineFunction (a,b,x):
    return a + b*x + 5*(random.random() - 1.0)
```

Generate and Plot Raw Data: (x,y) coordinates:

```
    5.00      3.96
    5.25      5.22
    5.50      4.30
    5.75      6.32

    ... data values removed ...

   19.00     14.44
   19.25     18.02
   19.50     18.62
   19.75     16.12
```

Principal Component Analysis

Two-dimensional plot of raw data:

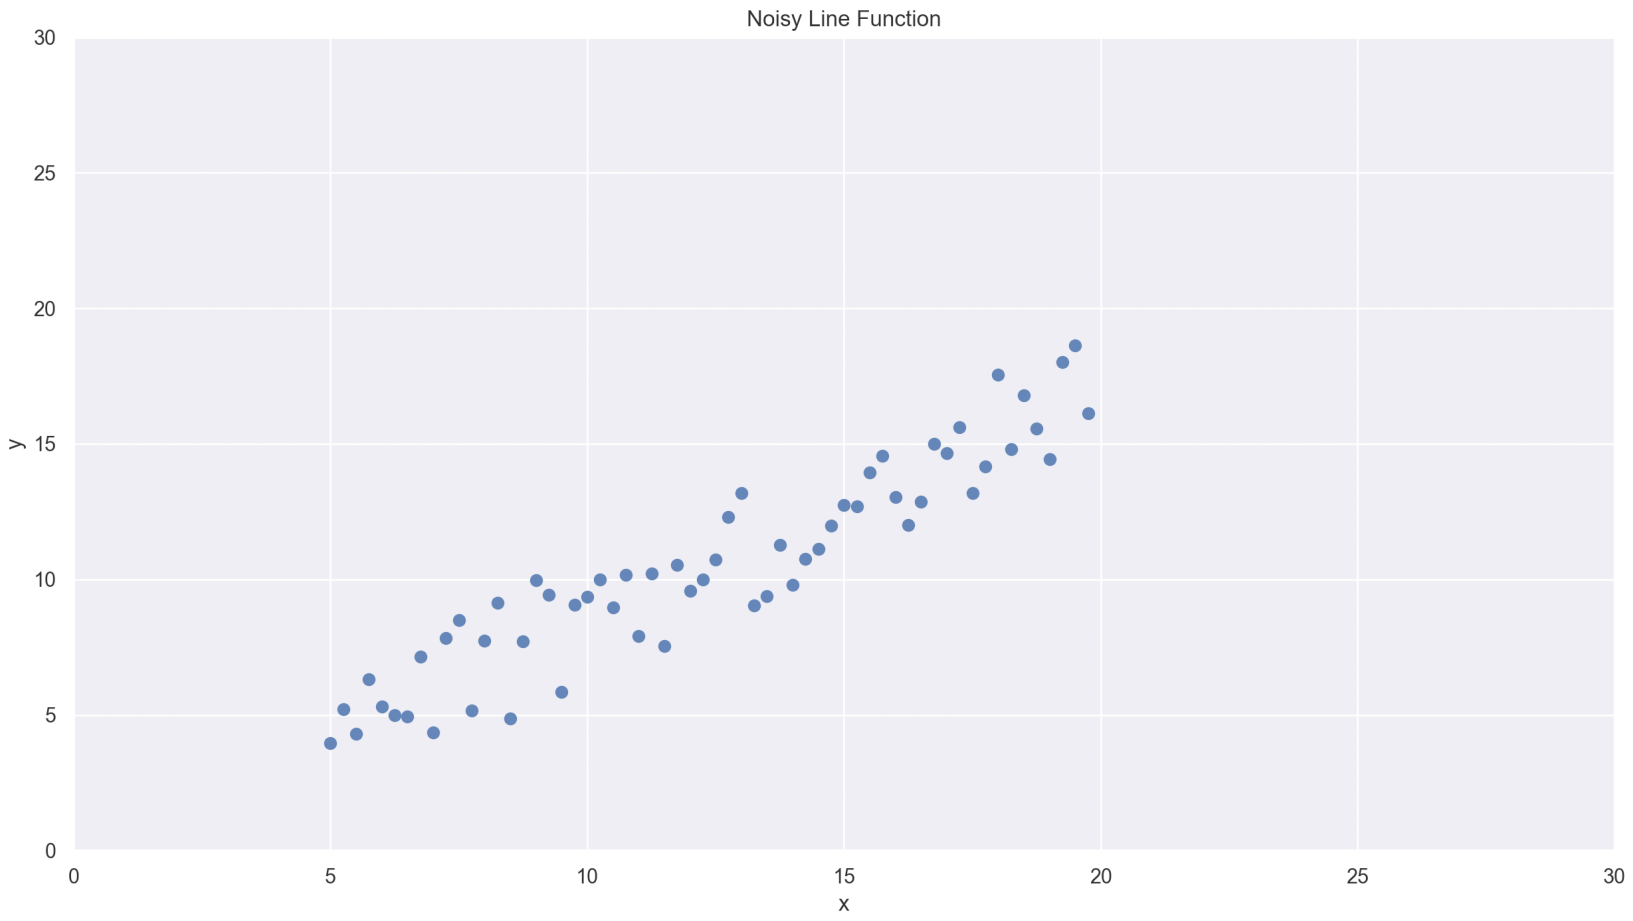

Principal Component Analysis

Compute Principal Components (No components = 2):

```
pca = PCA(n_components=2)
pca.fit(X)
```

Mean values; Eigenvalue and Eigenvectors

```
--- Print mean ...
[12.375     10.5313869]

--- Print components (first and second eigenvectors) ...

[ [ 0.75668904  0.65377496]     <-- first eigenvector ...
  [-0.65377496  0.75668904] ]   <-- second eigenvector ...

--- Print variance ...
[32.50694888  1.05218475]
```

Principal Component Analysis

Principal Components in Two Dimensions:

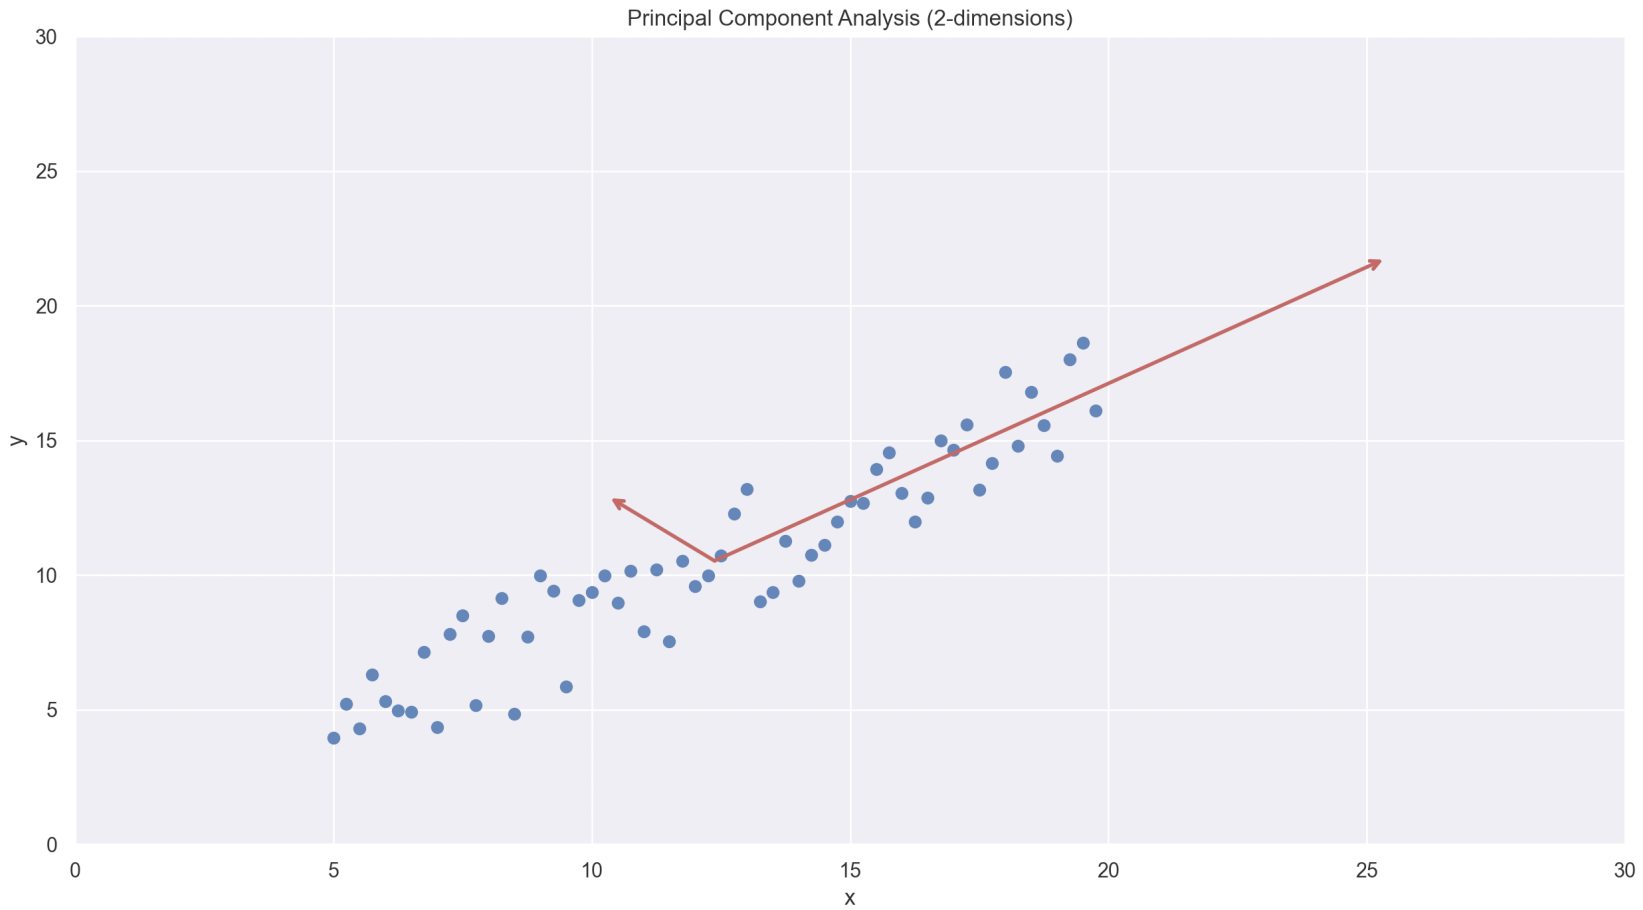

Principal Component Analysis

Dimensionality Reduction Platform (one dimension)

```
pca = PCA(n_components=1)
pca.fit(X)
X_pca = pca.transform(X)

print("--- original shape:   ", X.shape)
print("--- transformed shape:", X_pca.shape)

# Compute inverse transform on reduced data ...

X_new = pca.inverse_transform(X_pca)
```

Shape of Original and Transformed Data

```
--- original shape:    (60, 2)
--- transformed shape: (60, 1)
```

Principal Component Analysis

Principal Component Analysis in One Dimension:

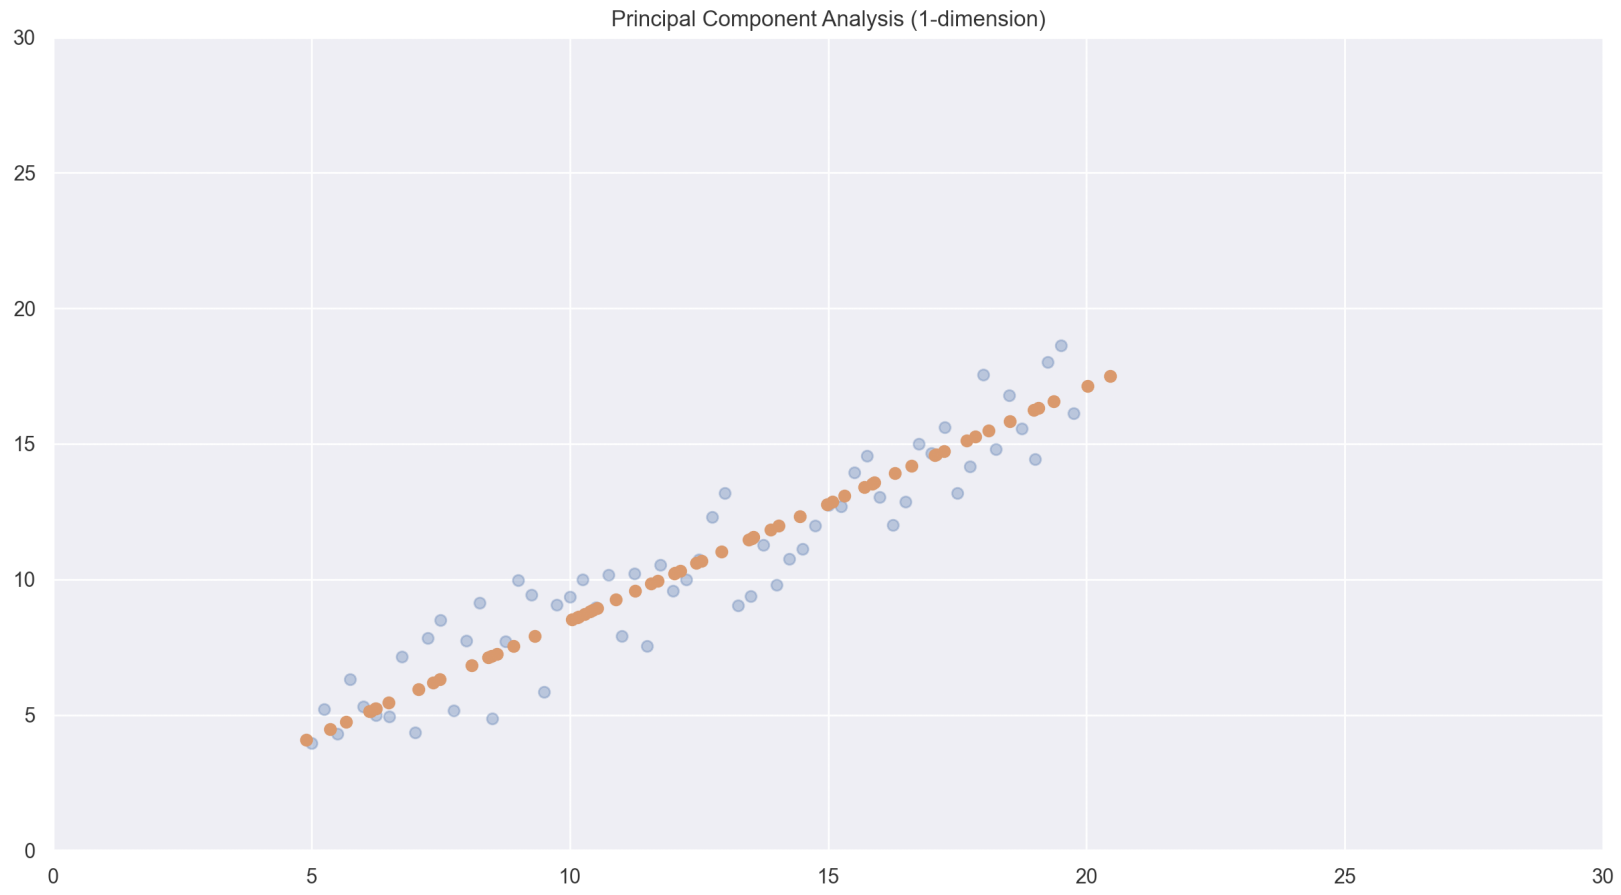

Principal Component Analysis

Components and Variance (as above)

```
--- Print components ...
[[0.75668904 0.65377496]]
```

```
--- Print variance ...
[32.50694888]
```

Side-by-Side Comparison of Coordinates

2D Coords (x,y)		--> 1D Coord System	--> Inverse Transform Coords	
5.00	3.96	-9.88	4.90	4.07
5.25	5.22	-8.87	5.67	4.74
5.50	4.30	-9.27	5.36	4.47
....	....	....	....	....
19.00	14.44	7.57	18.10	15.48
19.25	18.02	10.10	20.02	17.13
19.50	18.62	10.68	20.46	17.51
19.75	16.12	9.23	19.36	16.57

Principal Component Analysis

Strengths:

- Principal component analysis learns the linear transformation of data.
- Principal component analysis represents data in lower dimensions via an optimal orthogonal tranformation.
- Implementations are fast.

Weaknesses:

- Principal component analysis learns the linear transformation of data.
- As the number of features increases, the chances of overfitting of the model decreases significantly.

Source Code: See: python-code.d/autoencoder/

AutoEncoder vs PCA

Autoencoder vs PCA

A linearly activated autoencoder approximates principal component analysis. Mathematically, minimizing the reconstruction error in PCA modeling is the same as a single layer autoencoder.

Extensions of Autoencoder to Nonlinear Spaces

Autoencoders are nonlinear extensions of PCA.

Applications

Example 1. Simplest Example

Simplest Example. No nonlinear transformation.

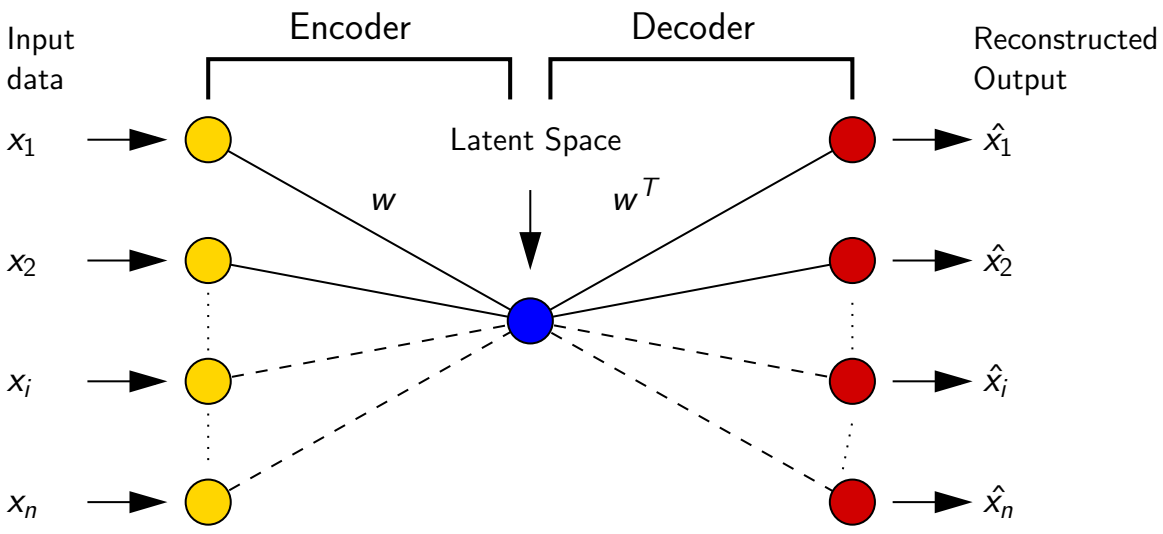

Features:

- A single hidden unit.
- Hidden unit has a **linear activation**.

Use same weight vector w for encoder/decoder. This is PCA.

Example 1. Simplest Example

Schematic for Tied Weights:

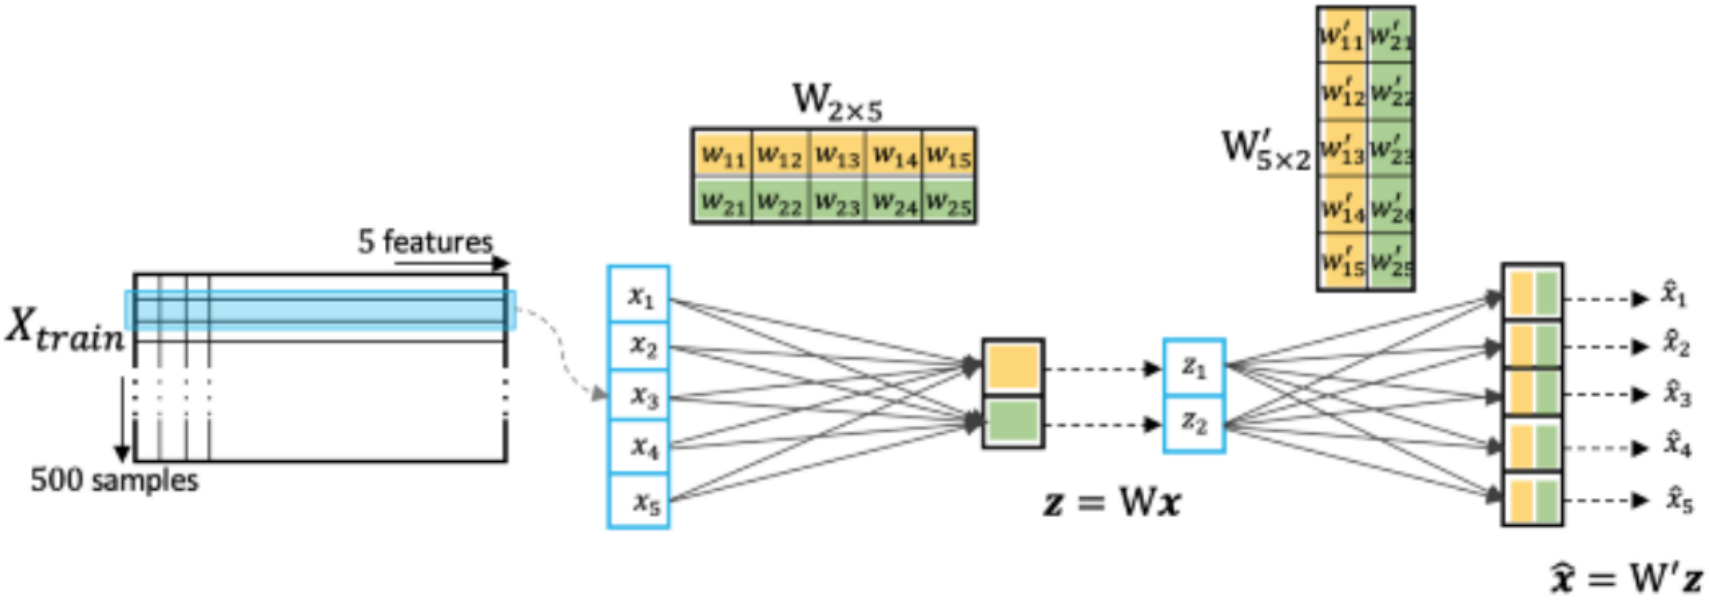

Required Constraints.

1. Tied weights, $W' = W^T$.
2. Orthogonal weights, $W^T W = I$.
3. Uncorrelated Encodings, $\mathrm{cor}(z_i, z_j) = 0, \text{if } i \neq j$.
4. Weights are Unit Norm, $\sum_{j=1}^{p} w_{ij}^2 = 1, i = 1, \ldots, k$.

Input Layer.
Size 5 x 1.

Encoding Layer.
Input 5 x 1.
Output 2 x 1.
Weights $W_{2\times5}$

Encodings,
output of the
Encoding Layer.
Size 2 x 1.

Decoding Layer.
Input 2 x 1.
Output 5 x 1.
Weights $W'_{5\times2}$

Example 1. Simplest Example

Schematic for Tied Weights (Autoencoder vs PCA):

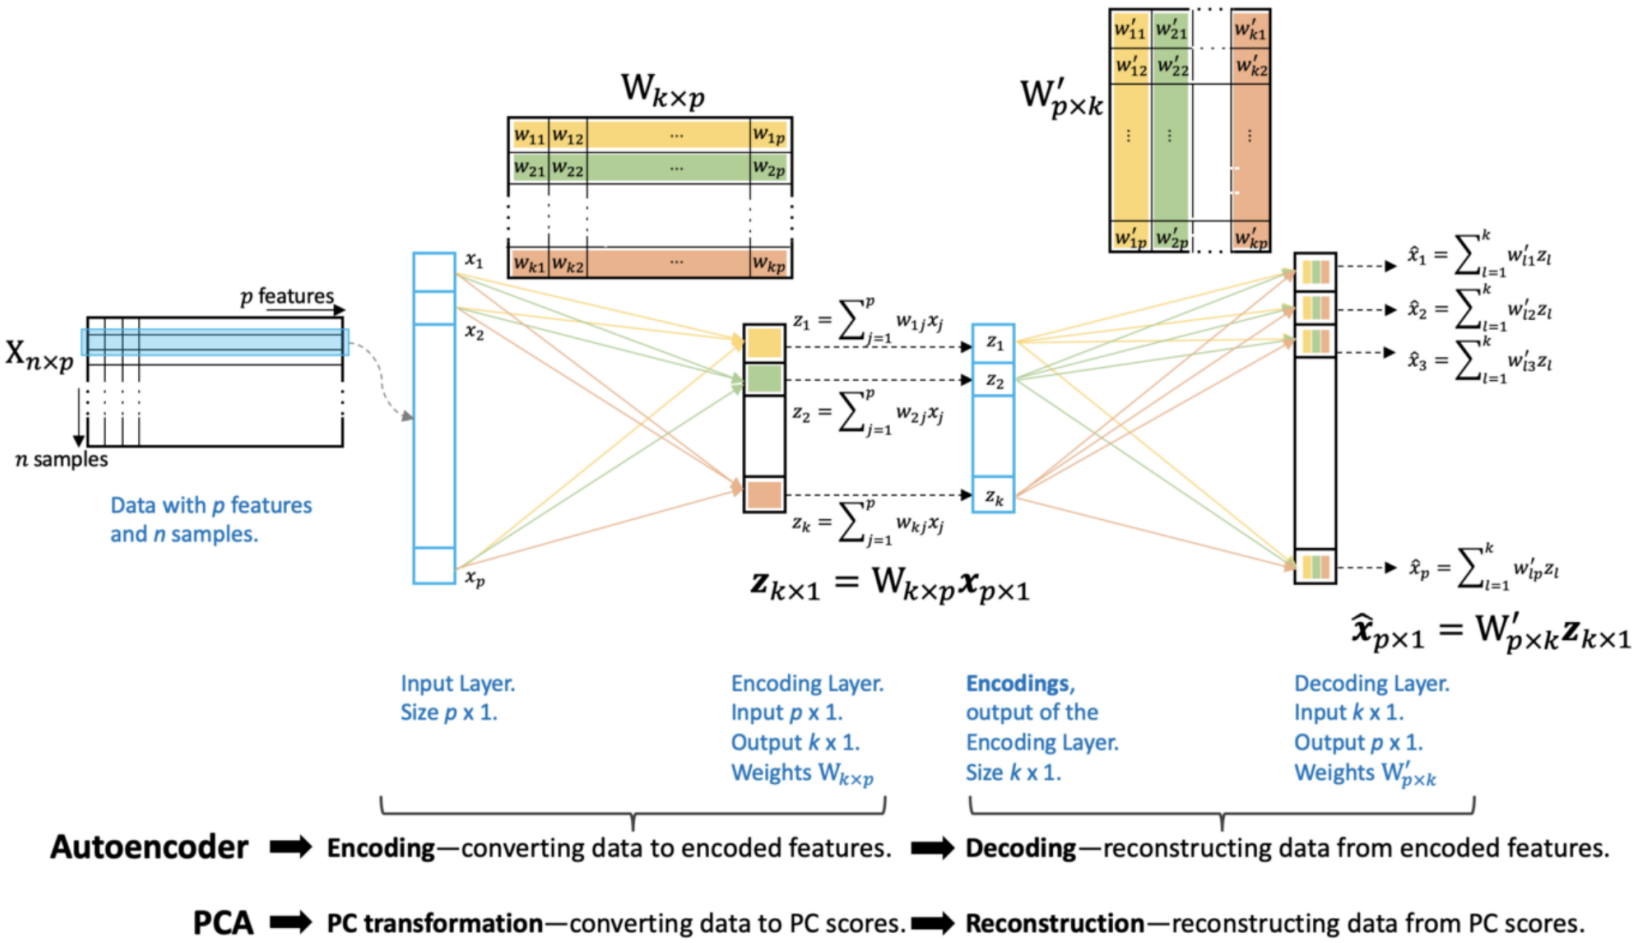

Example 1. Simplest Example

Theoretical Considerations:

Learn network weights by minimizing $L2$ divergence, i.e.,

$$L_2^2 = \underset{W}{arg\ min}\ E\left[\|x - \hat{x}\|^2\right] = \underset{W}{arg\ min}\ E\left[\|x - w^T wx\|^2\right] \tag{9}$$

Rewriting equation 9 as a matrix summation:

$$L_2^2 = \underset{W}{arg\ min} \sum_{i=1}^{n} \left[x_i - w^T wx_i\right]^T \left[x_i - w^T wx_i\right] \tag{10}$$

Here:

- x_i is a $(p \times 1)$ (p-dimensional) vector for the i-th data input.
- w is a $(1 \times p)$ vector of weights.

Example 1. Simplest Example

Training Procedure (Cont'd)"

$$L_2^2 = \arg\min_W \sum_{i=1}^{n} \left[x_i^T - x_i^T w^T w \right] \left[x_i - w^T w x_i \right].$$
$$= \arg\min_W \sum_{i=1}^{n} \left[x_i^T x_i - 2 x_i^T w^T w x_i + x_i^T w^T w w^T w x_i \right].$$

Equation 9 will be minimized when $w^T w$, a $(p \times p)$ matrix, equals I.

Example 1. Simplest Example

Generate Two-Dimensional Linear Dataset:

```
def StraightLineFunction (a,b,x):
    return a + b*x;

....

xvalues = list( np.arange( 5.0, 26.0, 1.00 ) );
yvalues = []
for x in xvalues:
    yvalue = StraightLineFunction( 0.0, 1.0, x)
    yvalues.append( yvalue );
```

Abbreviated Output:

```
[[ 5.  5.]                    [22. 22.]
 [ 6.  6.]                    [23. 23.]
 [ 7.  7.]                    [24. 24.]

 ........                     [25. 25.]]
                              (21, 2)
```

Example 1. Simplest Example

Statistics of Dataset:

```
--- Data mean: ...
[15. 15.]

--- Data covariance: ...
[[38.5 38.5]
 [38.5 38.5]]

--- Eigenvalues: ...
[77.  0.]

--- Eigenvectors: ...
[[ 0.70710678  0.70710678]
 [-0.70710678  0.70710678]]

--- Orthogonality of eigenvectors ...
[[1. 0.]
 [0. 1.]]
```

Example 1. Simplest Example

Principal Component Analysis:

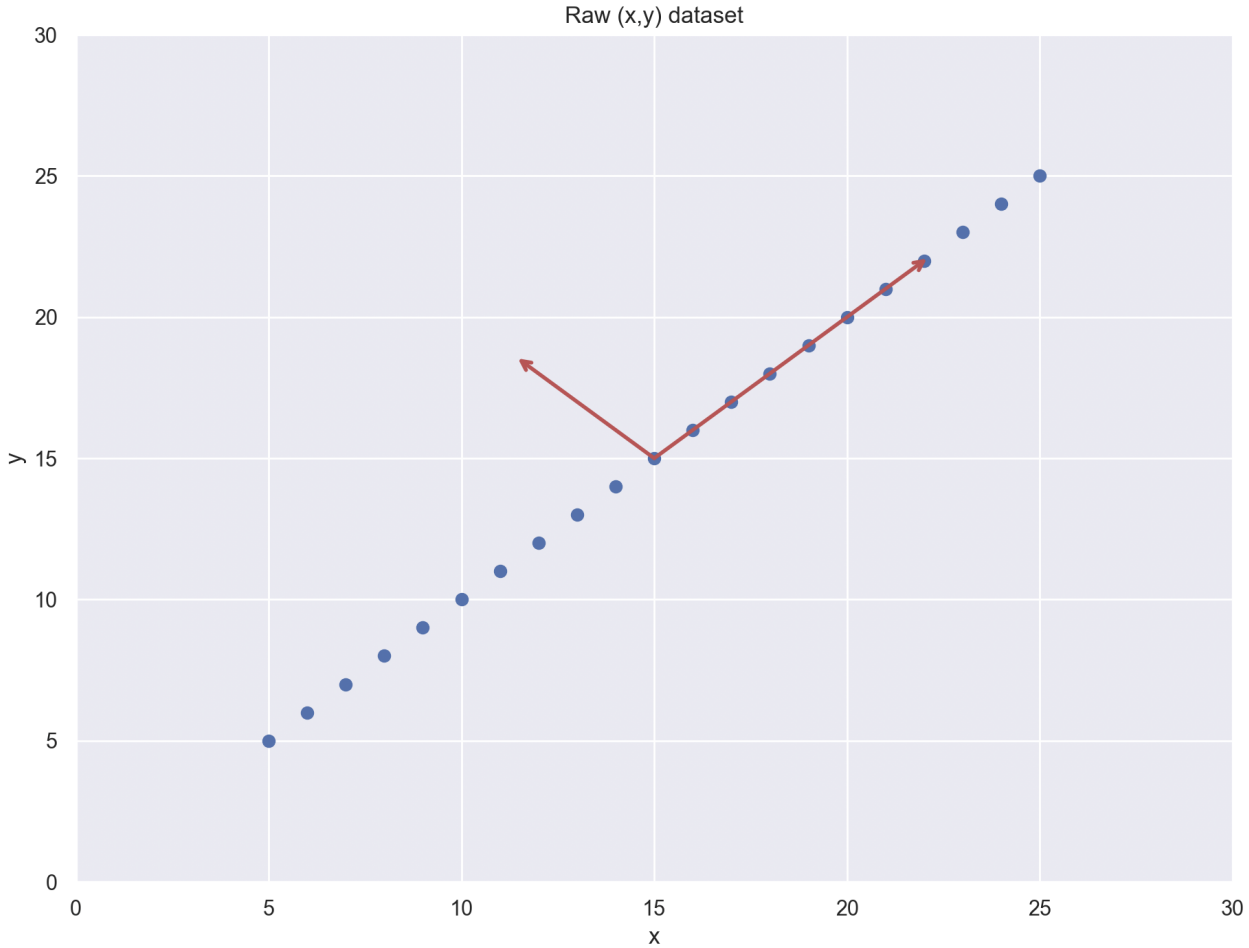

Example 1. Simplest Example

Create Train and Test Datasets. Scale Values:

```
--- X_train_scaled data ...
    [ [ 0.90  0.90 ]
      [ 0.15  0.15 ]
      [ 0.80  0.80 ]
      [ 0.40  0.40 ]
      [ 1.00  1.00 ]
      [ 0.75  0.75 ]
      [ 0.70  0.70 ]
      [ 0.60  0.60 ]
      [ 0.00  0.00 ]
      [ 0.05  0.05 ]
      [ 0.85  0.85 ]
      [ 0.45  0.45 ]
      [ 0.55  0.55 ]
      [ 0.50  0.50 ]
      [ 0.30  0.30 ]
      [ 0.95  0.95 ]
      [ 0.10  0.10 ]
      [ 0.65  0.65 ] ]
      (18, 2)
```

```
--- X_test_scaled data ...
    [ [  0.20  0.20 ]
      [  0.35  0.35 ]
      [  0.25  0.25 ] ]
      (3, 2)


--- Original test data ...
    [ [  9.00  9.00 ]
      [ 12.00 12.00 ]
      [ 10.00 10.00 ] ]
      (3, 2)
```

Example 1. Simplest Example

Principal Component Analysis: sklearn.decompostion.PCA()

```
pca = decomposition.PCA(n_components=2)
pca.fit( X_train_scaled )
```

Summary of PCA Results:

```
--- w_pca (components) ...

    [ [ 0.70710678  0.70710678]
      [-0.70710678  0.70710678] ]

--- w_pca * w_pca^T ...

    [ [ 1.00 -0.00 ]
      [-0.00  1.00 ] ]

--- Explained Variance Ratio:

    [ 1.000000e+00 9.966362e-35 ]
```

```
--- Fit model, apply transformation

    [ [-9.4280904e-02 -1.1573652e-18 ]
      [ 1.1785113e-01 -8.1142219e-19 ] ]

--- Data Reconstruction ...

    [ [ 0.20  0.20 ]
      [ 0.35  0.35 ]
      [ 0.25  0.25 ] ]

--- Reconstruction Loss: 0.0000
```

Example 1. Simplest Example

Autoencoder Architecture: Standard model:

Model: "sequential"

Layer (type)	Output Shape	Param #
dense (Dense)	(None, 1)	3
dense_1 (Dense)	(None, 2)	4

Total params: 7 (28.00 B)
Trainable params: 7 (28.00 B)
Non-trainable params: 0 (0.00 B)

```
encoder = Dense( encoding_dim, activation="linear", input_shape=(input_dim,),
                 use_bias = True)
decoder = Dense( input_dim, activation="linear", use_bias = True )

autoencoder = Sequential()
autoencoder.add(encoder)
autoencoder.add(decoder)

autoencoder.compile(metrics=['accuracy'], loss='mean_squared_error',
                    optimizer='sgd')
```

Example 1. Simplest Example

Summary: Linear activation functions.

Two layers:

- **Layer 1:** Input has two nodes, output is one node. Dense connectivity. Two weights + one bias.
- **Layer 2:** Input is one node; output is two nodes. Dense connectivity. Two weights + two bias parameters.

Example 1. Simplest Example

Time History of Learning:

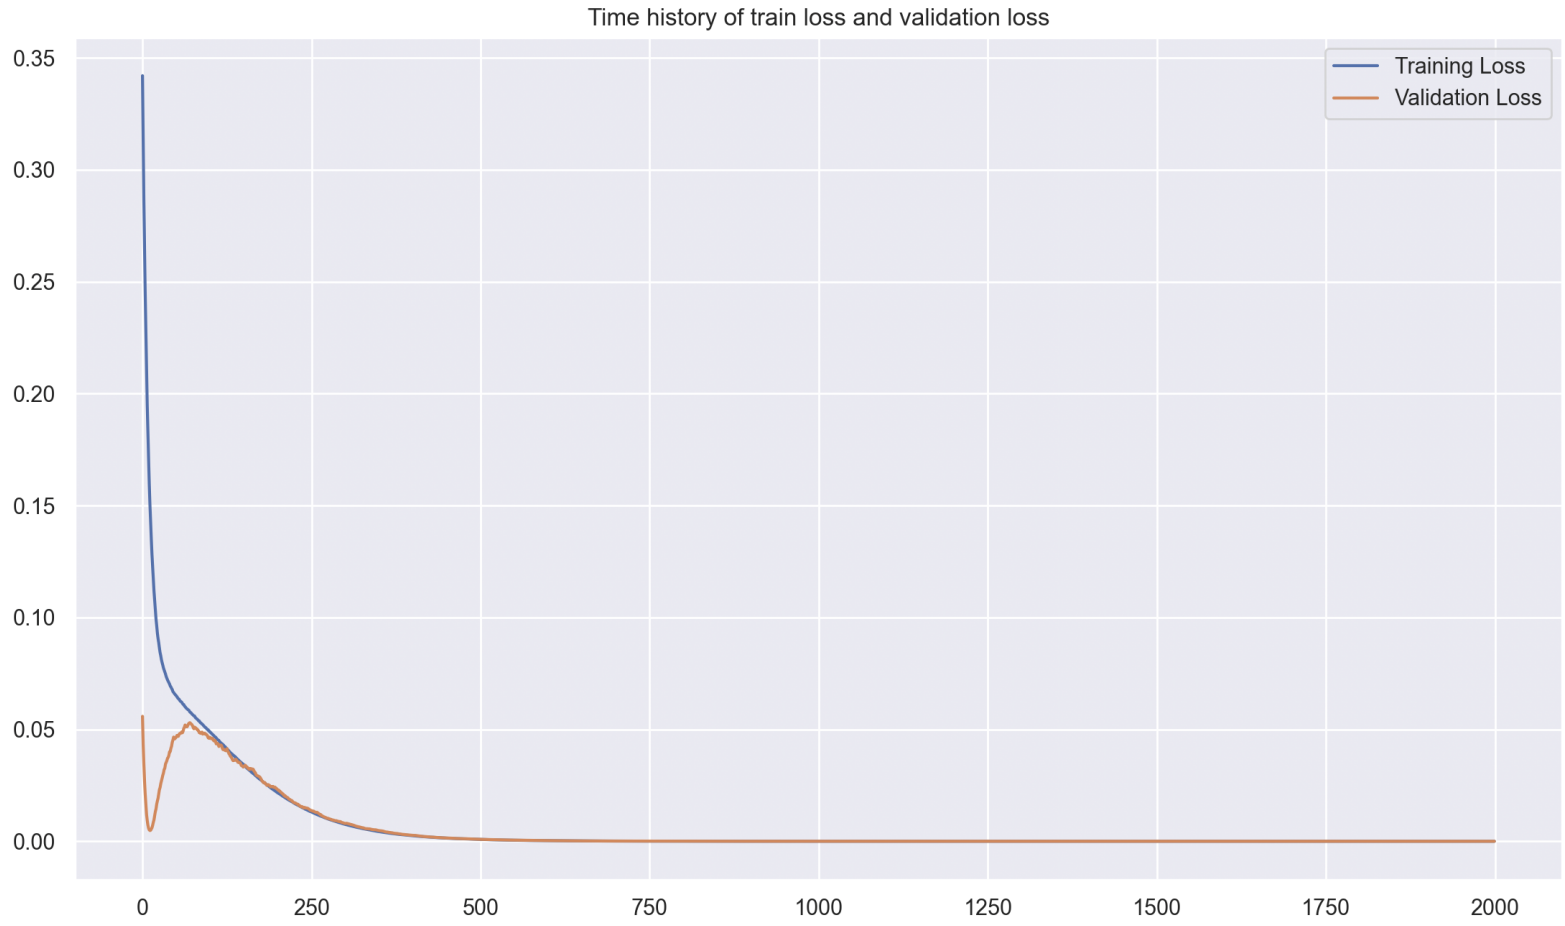

Example 1. Simplest Example

Training Results:

```
--- Encoder weights:
    [ [ 1.79  -0.6 ] ]

--- Encoder bias:
    [ -0.17 ]

--- Decoder weights:
    [ [ 0.84  0.84 ] ]

--- Decoder bias:
    [ 0.14  0.14 ]
```

Testing Results:

```
--- X_test_scaled:   --->
[[0.2  0.2 ]
 [0.35 0.35]
 [0.25 0.25]]

--- Autoencoder Predictions:
    [ [ 0.20  0.20 ]
      [ 0.35  0.35 ]
      [ 0.25  0.25 ] ]

--- Original Test Data: --->
[ [  9.  9.]
  [ 12. 12.]
  [ 10. 10.] ]

--- Autoencoder Predictions:
    [ [  9.   9. ]
      [ 12.  12. ]
      [ 10.  10. ] ]
```

Example 1. Simplest Example

Autoencoder Architecture: Tied Weights. No bias parameters.

Model: "sequential_1"

Layer (type)	Output Shape	Param #
dense_2 (Dense)	(None, 1)	3
dense_tied (DenseTied)	(None, 2)	3

Total params: 3 (12.00 B)
Trainable params: 3 (12.00 B)
Non-trainable params: 0 (0.00 B)

Example 1. Simplest Example

Time History of Learning:

Example 1. Simplest Example

Training Results: Tied Weights

```
--- Encoder weights:
    [ [ 0.305  0.878 ] ]

--- Encoder bias:
    [ None ]

--- Decoder weights:
    [ [ 0.305  0.878 ] ]

--- Decoder bias:
    [ None ]
```

Testing Results: Not great ...

```
--- X_test_scaled:   --->
[[0.2  0.2 ]
 [0.35 0.35]
 [0.25 0.25]]

--- Autoencoder Predictions:
[ [ 0.11  0.24 ]
  [ 0.19  0.42 ]
  [ 0.13  0.30 ] ]

--- Original Test Data: --->
[ [  9.  9.]
  [ 12. 12.]
  [ 10. 10.] ]

--- Autoencoder Predictions:
[ [ 4.77  10.86 ]
  [ 6.36  14.48 ]
  [ 5.30  12.07 ] ]
```

Example 1. Simplest Example

Autoencoder Architecture: Tied Weights. Bias parameters.

Model: "sequential_2"

Layer (type)	Output Shape	Param #
dense_3 (Dense)	(None, 1)	3
dense_tied_1 (DenseTied)	(None, 2)	5

Total params: 5 (20.00 B)
Trainable params: 5 (20.00 B)
Non-trainable params: 0 (0.00 B)

Example 1. Simplest Example

Time History of Learning:

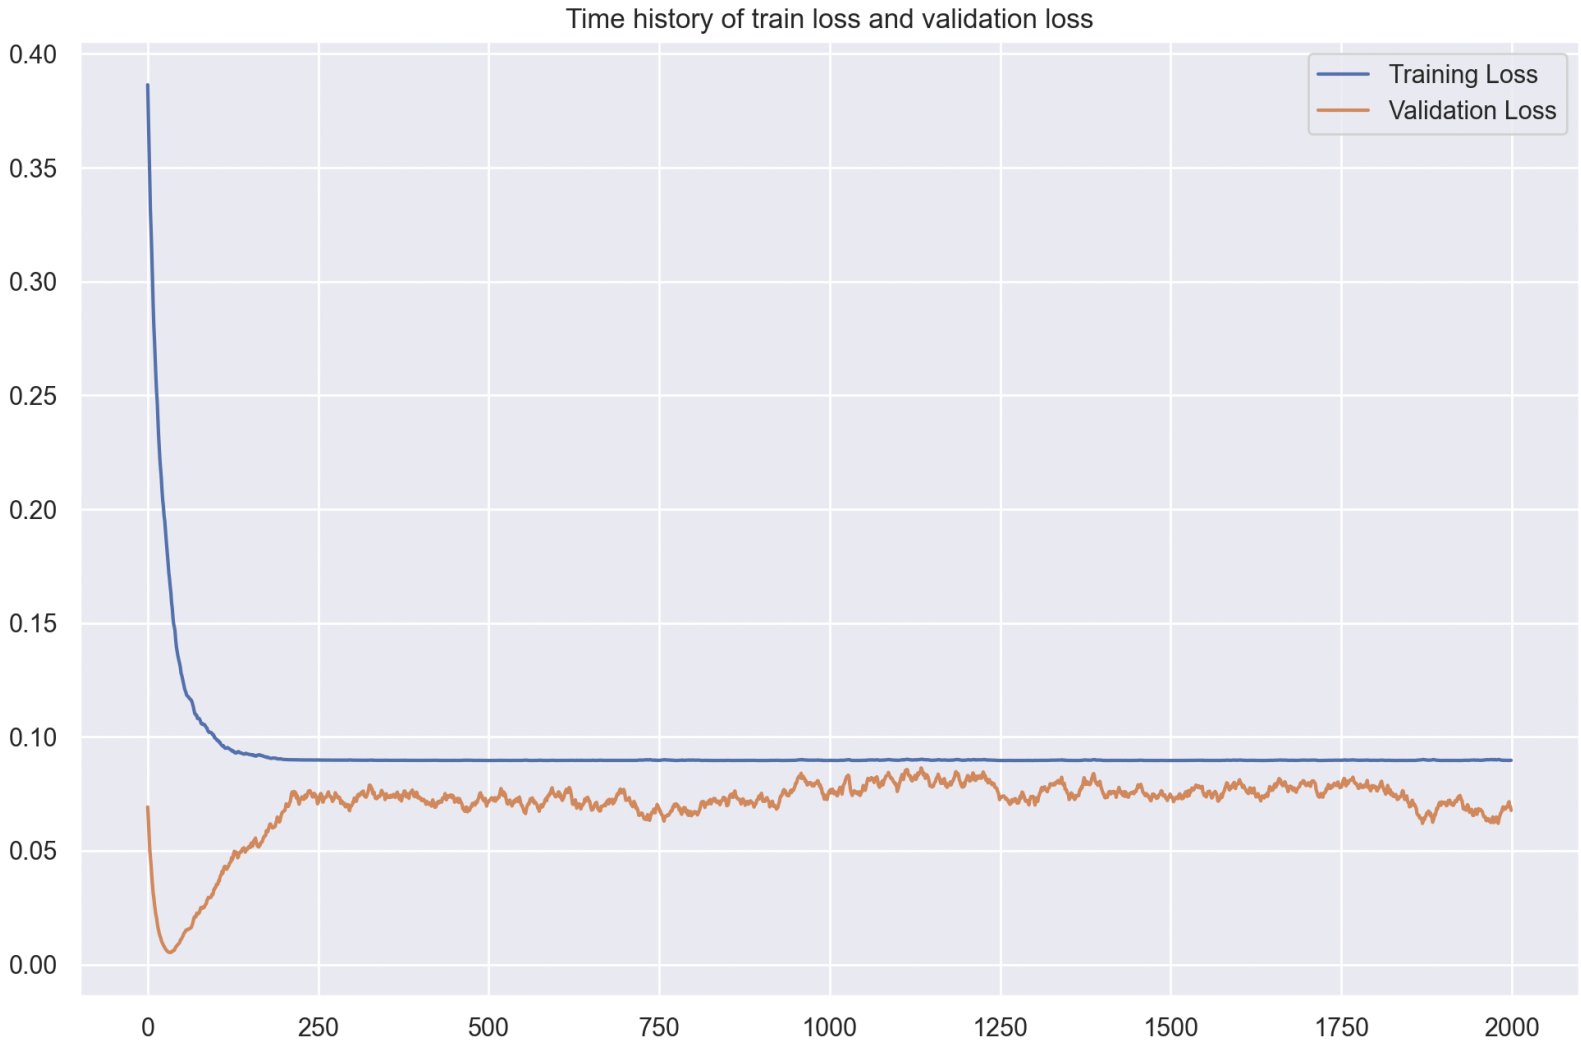

Example 1. Simplest Example

Training Results:

```
--- Encoder weights:
    [ [ 1.098 -0.847 ] ]

--- Encoder bias:
    [  0.00 ]

--- Decoder weights:
    [ [ 0.84  0.84 ] ]

--- Decoder bias:
    [ 0.373  0.635 ]
```

Testing Results: Inferior performance ...

```
--- X_test_scaled:   --->
[[0.2  0.2 ]
 [0.35 0.35]
 [0.25 0.25]]

--- Autoencoder Predictions:
    [ [ 0.43  0.60 ]
      [ 0.48  0.57 ]
      [ 0.45  0.59 ] ]

--- Original Test Data: --->
[ [  9.  9.]
  [ 12. 12.]
  [ 10. 10.] ]

--- Autoencoder Predictions:
    [ [ 3.14   -0.99 ]
      [ 4.06   -1.52 ]
      [ 3.45   -1.17 ] ]
```

Example 2. Linear vs Nonlinear Dimensionality Reduction

Linear vs Nonlinear Dimensionality Reduction:

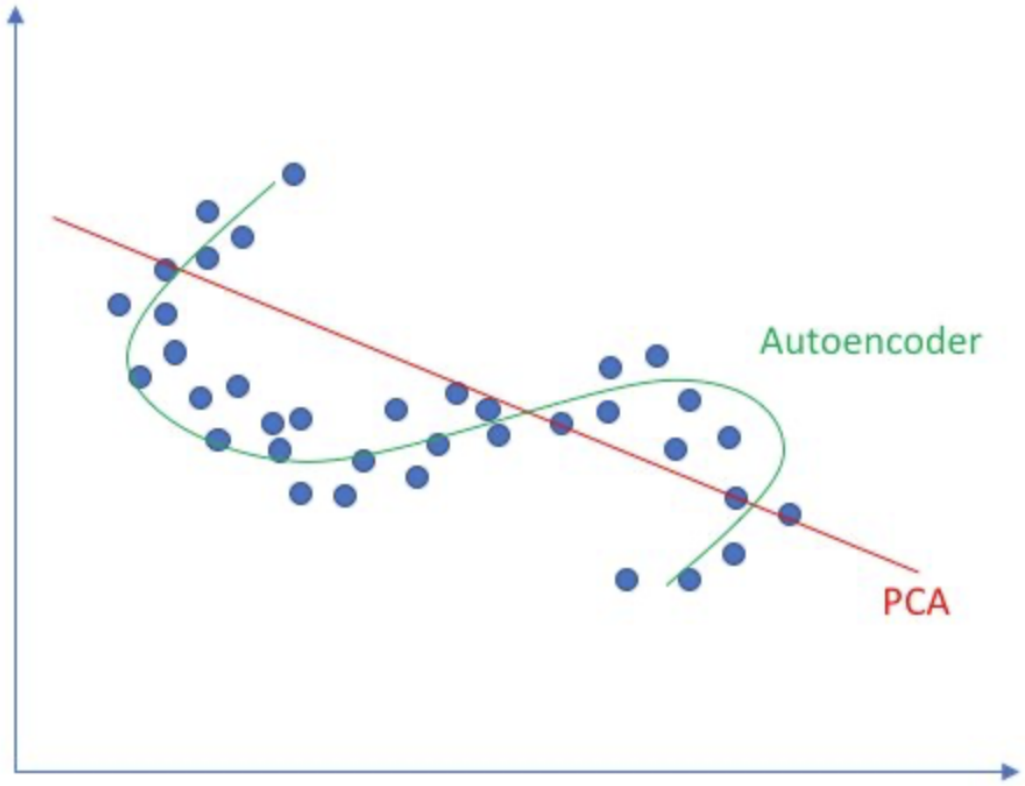

Source: Nugroho H. et al., 2020

Example 2. Linear vs Nonlinear Dimensionality Reduction

Linear vs Nonlinear Dimensionality Reduction:

Function	Feature Space	PCA Reconstruction	Auto Encoder Reconstruction
$y=mx+c$	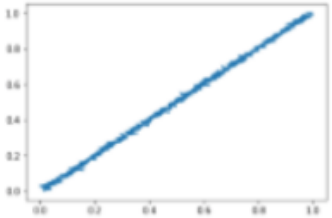	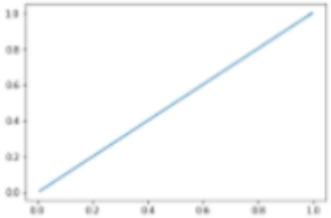	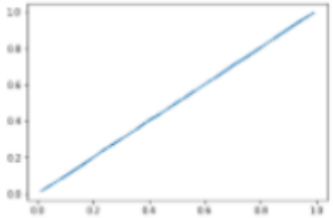
$y=mx^2+c$	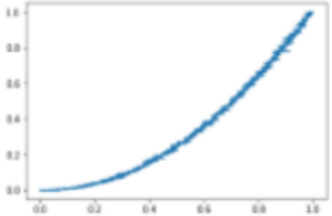	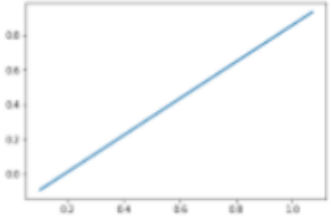	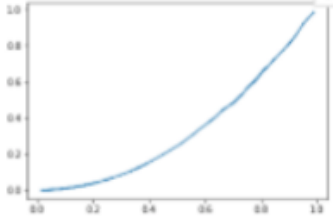
$y=mx^8+c$	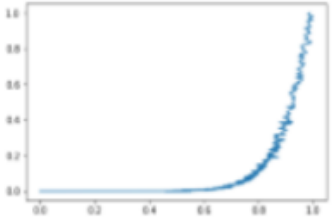	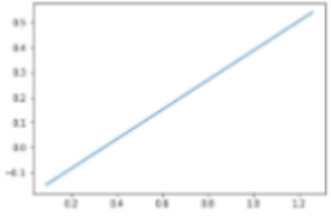	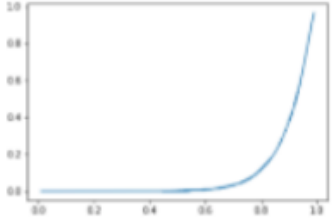

Example 2. Linear vs Nonlinear Dimensionality Reduction

DL4J: Dataset

Example 2. Linear vs Nonlinear Dimensionality Reduction

DL4J: Neural Network Architecture

Example 3. Anomaly Detection

Example 4. Handwritten Digit Recognition

Problem Statement

- Demonstrate anomaly detection on MNIST using simple autoencoder.
- Goal is to identify digits that are unusual.

MNIST Handwritten Digit Dataset

- MNIST (Modified National Institute of Standards and Technology) is a database of handwritten digits provided by NIST.
- The database contains: 60,000 training images and 10,000 testing images.
- Each image is a scan of a handwritten image 0 through 9.
- Differences among images are due to variations in handwriting style.

Example 4. Handwritten Digit Recognition

Sample Digits

Example 4. Handwritten Digit Recognition

Solution Procedure

- Images are centered on a 28x28 grid of pixels. Individual pixels take a value 0 through 255.
- Individual 28x28 images $\longrightarrow$ 1x784 data vector.

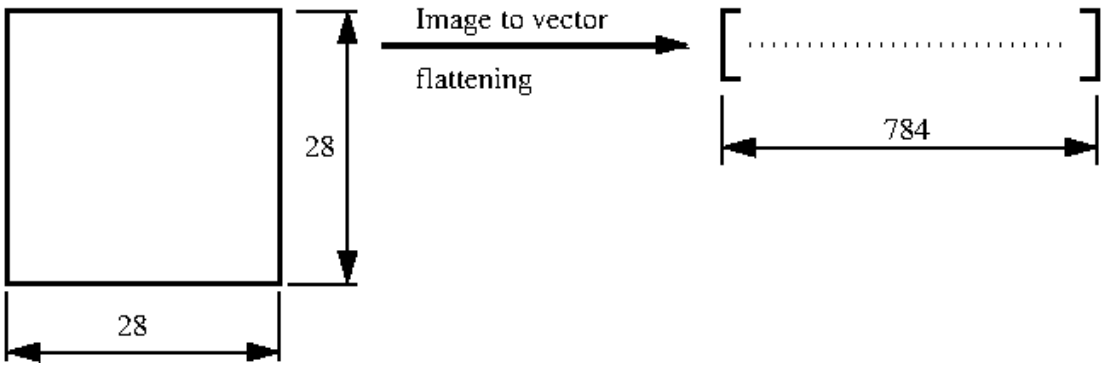

- Create network configutation with 784 inputs/outputs, contracting down to a ten-dimensional embedding vector, i.e., 784 -> 250 -> 10 -> 250 -> 784.

Example 4. Handwritten Digit Recognition

DL4J: Training Dataset

DL4J: Testing Dataset

Example 4. Handwritten Digit Recognition

DL4J: Network Configuration: 784 -> 250 -> 10 -> 250 -> 784.

```
MultiLayerConfiguration conf = new NeuralNetConfiguration.Builder()
        .seed(12345)
        .weightInit(WeightInit.XAVIER)
        .updater(new AdaGrad(0.05))
        .activation(Activation.RELU)
        .l2(0.0001)
        .list()
        .layer(new DenseLayer.Builder().nIn(784).nOut(250)
            .build())
        .layer(new DenseLayer.Builder().nIn(250).nOut(10)
            .build())
        .layer(new DenseLayer.Builder().nIn(10).nOut(250)
            .build())
        .layer(new OutputLayer.Builder().nIn(250).nOut(784)
            .activation(Activation.LEAKYRELU)
            .lossFunction(LossFunctions.LossFunction.MSE)
            .build())
        .build();

MultiLayerNetwork net = new MultiLayerNetwork(conf);
net.setListeners(Collections.singletonList(new ScoreIterationListener(10)));
```

Example 4. Handwritten Digit Recognition

DL4J: Summary of Network Model.

```
==========================================================================
LayerName (LayerType)    nIn,nOut    Params   ParamsShape
==========================================================================
layer0    (DenseLayer)    784,250   196,250   W:{784,250}, b:{1,250}
layer1    (DenseLayer)     250,10     2,510   W:{250,10}, b:{1,10}
layer2    (DenseLayer)     10,250     2,750   W:{10,250}, b:{1,250}
layer3    (OutputLayer)   250,784   196,784   W:{250,784}, b:{1,784}
--------------------------------------------------------------------------
Total Parameters: 398,294  Trainable Parameters: 398,294
==========================================================================
```

Example 4. Handwritten Digit Recognition

Training Procedure:

Compute Reconstruction Error:

Example 4. Handwritten Digit Recognition

Results: Best (left) and Worst (right) Results

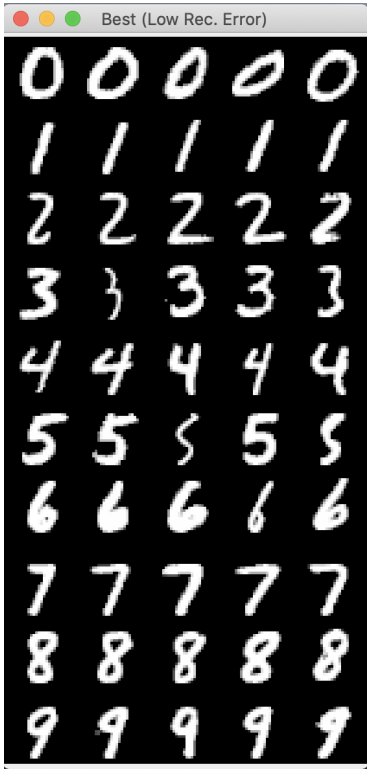

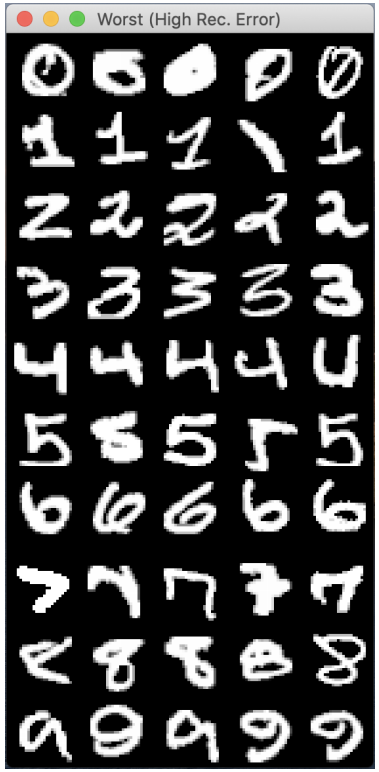

[Best, Worst] = [low, high] reconstruction error.

Example 4. Handwritten Digit Recognition

Strengths

- ...
- ...
- ...

Weaknesses

- Flattening the image looses spatial information, which, in turn, can simplify the learning process.
- ...

References

- Aggarwal C.C., Neural Networks and Deep Learning, Springer, 2018.
- Nugroho H., Susanty M., Irawan A., Koyimatu M., and Yunita A., Fully Convolutional Variational Autoencoder for Feature Extraction of Fire Detection System, Journal of Computer Science and Engineering, Vol. 13, No. 1, 2020.
- Watt J., Borhani R., Katsaggelos A.K., Machine Learning Refined, Second Edition, Cambridge University Press, 2020.

Appendix A

Cheatsheets