

Chapter 8 – Dimensionality Reduction

This notebook contains all the sample code and solutions to the exercises in chapter 8.



[Run in Google Colab \(https://colab.research.google.com/github/ageron/handson-ml2/blob/master/08_dimensionality_reduction.ipynb\)](https://colab.research.google.com/github/ageron/handson-ml2/blob/master/08_dimensionality_reduction.ipynb)

Setup

First, let's import a few common modules, ensure Matplotlib plots figures inline and prepare a function to save the figures. We also check that Python 3.5 or later is installed (although Python 2.x may work, it is deprecated so we strongly recommend you use Python 3 instead), as well as Scikit-Learn ≥ 0.20 .

```
In [1]: # Python  $\geq 3.5$  is required
import sys
assert sys.version_info >= (3, 5)

# Scikit-Learn  $\geq 0.20$  is required
import sklearn
assert sklearn.__version__ >= "0.20"

# Common imports
import numpy as np
import os

# to make this notebook's output stable across runs
np.random.seed(42)

# To plot pretty figures
%matplotlib inline
import matplotlib as mpl
import matplotlib.pyplot as plt
mpl.rc('axes', labelsize=14)
mpl.rc('xtick', labelsize=12)
mpl.rc('ytick', labelsize=12)

# Where to save the figures
PROJECT_ROOT_DIR = "."
CHAPTER_ID = "dim_reduction"
IMAGES_PATH = os.path.join(PROJECT_ROOT_DIR, "images", CHAPTER_ID)
os.makedirs(IMAGES_PATH, exist_ok=True)

def save_fig(fig_id, tight_layout=True, fig_extension="png", resolution=300):
    path = os.path.join(IMAGES_PATH, fig_id + "." + fig_extension)
    print("Saving figure", fig_id)
    if tight_layout:
        plt.tight_layout()
    plt.savefig(path, format=fig_extension, dpi=resolution)

# Ignore useless warnings (see SciPy issue #5998)
import warnings
warnings.filterwarnings(action="ignore", message="^internal gelsd")
```

Projection methods

Build 3D dataset:

```
In [2]: np.random.seed(4)
m = 60
w1, w2 = 0.1, 0.3
noise = 0.1

angles = np.random.rand(m) * 3 * np.pi / 2 - 0.5
X = np.empty((m, 3))
X[:, 0] = np.cos(angles) + np.sin(angles)/2 + noise * np.random.randn(m) / 2
X[:, 1] = np.sin(angles) * 0.7 + noise * np.random.randn(m) / 2
X[:, 2] = X[:, 0] * w1 + X[:, 1] * w2 + noise * np.random.randn(m)
```

PCA using Scikit-Learn

With Scikit-Learn, PCA is really trivial. It even takes care of mean centering for you:

```
In [3]: from sklearn.decomposition import PCA

pca = PCA(n_components = 2)
X2D = pca.fit_transform(X)
```

```
In [4]: X2D[:5]
```

```
Out[4]: array([[ 1.26203346,  0.42067648],
               [-0.08001485, -0.35272239],
               [ 1.17545763,  0.36085729],
               [ 0.89305601, -0.30862856],
               [ 0.73016287, -0.25404049]])
```

Notice that running PCA multiple times on slightly different datasets may result in different results. In general the only difference is that some axes may be flipped.

Recover the 3D points projected on the plane (PCA 2D subspace).

```
In [5]: X3D_inv = pca.inverse_transform(X2D)
```

Of course, there was some loss of information during the projection step, so the recovered 3D points are not exactly equal to the original 3D points:

```
In [6]: np.allclose(X3D_inv, X)
```

```
Out[6]: False
```

We can compute the reconstruction error:

```
In [7]: np.mean(np.sum(np.square(X3D_inv - X), axis=1))
```

```
Out[7]: 0.010170337792848549
```

The PCA object gives access to the principal components that it computed:

```
In [8]: pca.components_
Out[8]: array([[ -0.93636116, -0.29854881, -0.18465208],
               [ 0.34027485, -0.90119108, -0.2684542 ]])
```

Notice how the axes are flipped.

Now let's look at the explained variance ratio:

```
In [9]: pca.explained_variance_ratio_
Out[9]: array([0.84248607, 0.14631839])
```

The first dimension explains 84.2% of the variance, while the second explains 14.6%.

By projecting down to 2D, we lost about 1.1% of the variance:

```
In [10]: 1 - pca.explained_variance_ratio_.sum()
Out[10]: 0.011195535570688975
```

Next, let's generate some nice figures! :)

Utility class to draw 3D arrows (copied from <http://stackoverflow.com/questions/11140163> (<http://stackoverflow.com/questions/11140163>))

```
In [11]: from matplotlib.patches import FancyArrowPatch
         from mpl_toolkits.mplot3d import proj3d

         class Arrow3D(FancyArrowPatch):
             def __init__(self, xs, ys, zs, *args, **kwargs):
                 FancyArrowPatch.__init__(self, (0,0), (0,0), *args, **kwargs)
                 self._verts3d = xs, ys, zs

             def draw(self, renderer):
                 xs3d, ys3d, zs3d = self._verts3d
                 xs, ys, zs = proj3d.proj_transform(xs3d, ys3d, zs3d, renderer.M)
                 self.set_positions((xs[0],ys[0]),(xs[1],ys[1]))
                 FancyArrowPatch.draw(self, renderer)
```

Express the plane as a function of x and y.

```
In [12]: axes = [-1.8, 1.8, -1.3, 1.3, -1.0, 1.0]

         x1s = np.linspace(axes[0], axes[1], 10)
         x2s = np.linspace(axes[2], axes[3], 10)
         x1, x2 = np.meshgrid(x1s, x2s)

         C = pca.components_
         R = C.T.dot(C)
         z = (R[0, 2] * x1 + R[1, 2] * x2) / (1 - R[2, 2])
```

Plot the 3D dataset, the plane and the projections on that plane.

```

In [13]: from mpl_toolkits.mplot3d import Axes3D

fig = plt.figure(figsize=(6, 3.8))
ax = fig.add_subplot(111, projection='3d')

X3D_above = X[X[:, 2] > X3D_inv[:, 2]]
X3D_below = X[X[:, 2] <= X3D_inv[:, 2]]

ax.plot(X3D_below[:, 0], X3D_below[:, 1], X3D_below[:, 2], "bo", alpha=0.5)

ax.plot_surface(x1, x2, z, alpha=0.2, color="k")
np.linalg.norm(C, axis=0)
ax.add_artist(Arrow3D([0, C[0, 0]], [0, C[0, 1]], [0, C[0, 2]], mutation_scale
=15, lw=1, arrowstyle="-|>", color="k"))
ax.add_artist(Arrow3D([0, C[1, 0]], [0, C[1, 1]], [0, C[1, 2]], mutation_scale
=15, lw=1, arrowstyle="-|>", color="k"))
ax.plot([0], [0], [0], "k.")

for i in range(m):
    if X[i, 2] > X3D_inv[i, 2]:
        ax.plot([X[i][0], X3D_inv[i][0]], [X[i][1], X3D_inv[i][1]], [X
[i][2], X3D_inv[i][2]], "k-")
    else:
        ax.plot([X[i][0], X3D_inv[i][0]], [X[i][1], X3D_inv[i][1]], [X
[i][2], X3D_inv[i][2]], "k-", color="#505050")


ax.plot(X3D_inv[:, 0], X3D_inv[:, 1], X3D_inv[:, 2], "k+")
ax.plot(X3D_inv[:, 0], X3D_inv[:, 1], X3D_inv[:, 2], "k.")
ax.plot(X3D_above[:, 0], X3D_above[:, 1], X3D_above[:, 2], "bo")
ax.set_xlabel("$x_1$", fontsize=18, labelpad=10)
ax.set_ylabel("$x_2$", fontsize=18, labelpad=10)
ax.set_zlabel("$x_3$", fontsize=18, labelpad=10)
ax.set_xlim(axes[0:2])
ax.set_ylim(axes[2:4])
ax.set_zlim(axes[4:6])

# Note: If you are using Matplotlib 3.0.0, it has a bug and does not
# display 3D graphs properly.
# See https://github.com/matplotlib/matplotlib/issues/12239
# You should upgrade to a later version. If you cannot, then you can
# use the following workaround before displaying each 3D graph:
# for spine in ax.spines.values():
#     spine.set_visible(False)

save_fig("dataset_3d_plot")
plt.show()

```


Saving figure dataset_3d_plot



```
In [14]: fig = plt.figure()
ax = fig.add_subplot(111, aspect='equal')

ax.plot(X2D[:, 0], X2D[:, 1], "k+")
ax.plot(X2D[:, 0], X2D[:, 1], "k.")
ax.plot([0], [0], "ko")
ax.arrow(0, 0, 0, 1, head_width=0.05, length_includes_head=True, head_length=0.1, fc='k', ec='k')
ax.arrow(0, 0, 1, 0, head_width=0.05, length_includes_head=True, head_length=0.1, fc='k', ec='k')
ax.set_xlabel("$z_1$", fontsize=18)
ax.set_ylabel("$z_2$", fontsize=18, rotation=0)
ax.axis([-1.5, 1.3, -1.2, 1.2])
ax.grid(True)
save_fig("dataset_2d_plot")
```

Saving figure dataset_2d_plot



PCA Example : MNIST compression

```
In [15]: from sklearn.datasets import fetch_openml

mnist = fetch_openml('mnist_784', version=1)
mnist.target = mnist.target.astype(np.uint8)
```

```
In [16]: from sklearn.model_selection import train_test_split

X = mnist["data"]
y = mnist["target"]

X_train, X_test, y_train, y_test = train_test_split(X, y)
```


```
In [17]: pca = PCA()
pca.fit(X_train)
cumsum = np.cumsum(pca.explained_variance_ratio_)
d = np.argmax(cumsum >= 0.95) + 1
```

```
In [18]: d
```

```
Out[18]: 154
```

```
In [19]: plt.figure(figsize=(6,4))
plt.plot(cumsum, linewidth=3)
plt.axis([0, 400, 0, 1])
plt.xlabel("Dimensions")
plt.ylabel("Explained Variance")
plt.plot([d, d], [0, 0.95], "k:")
plt.plot([0, d], [0.95, 0.95], "k:")
plt.plot(d, 0.95, "ko")
plt.annotate("Elbow", xy=(65, 0.85), xytext=(70, 0.7),
            arrowprops=dict(arrowstyle="->"), fontsize=16)
plt.grid(True)
save_fig("explained_variance_plot")
plt.show()
```

Saving figure explained_variance_plot



```
In [20]: pca = PCA(n_components=0.95)
X_reduced = pca.fit_transform(X_train)
```

```
In [21]: pca.n_components_
```

```
Out[21]: 154
```

```
In [22]: np.sum(pca.explained_variance_ratio_)
```

```
Out[22]: 0.9503684424557437
```


```
In [23]: pca = PCA(n_components = 154)
X_reduced = pca.fit_transform(X_train)
X_recovered = pca.inverse_transform(X_reduced)
```

```
In [24]: def plot_digits(instances, images_per_row=5, **options):
size = 28
images_per_row = min(len(instances), images_per_row)
images = [instance.reshape(size,size) for instance in instances]
n_rows = (len(instances) - 1) // images_per_row + 1
row_images = []
n_empty = n_rows * images_per_row - len(instances)
images.append(np.zeros((size, size * n_empty)))
for row in range(n_rows):
    rimages = images[row * images_per_row : (row + 1) * images_per_row]
    row_images.append(np.concatenate(rimages, axis=1))
image = np.concatenate(row_images, axis=0)
plt.imshow(image, cmap = mpl.cm.binary, **options)
plt.axis("off")
```

```
In [25]: plt.figure(figsize=(7, 4))
plt.subplot(121)
plot_digits(X_train[:, :2100])
plt.title("Original", fontsize=16)
plt.subplot(122)
plot_digits(X_recovered[:, :2100])
plt.title("Compressed", fontsize=16)

save_fig("mnist_compression_plot")
```

Saving figure mnist_compression_plot



```
In [26]: X_reduced_pca = X_reduced
```

Kernel PCA

```
In [27]: from sklearn.datasets import make_swiss_roll
X, t = make_swiss_roll(n_samples=1000, noise=0.2, random_state=42)
```

```
In [28]: from sklearn.decomposition import KernelPCA

rbf_pca = KernelPCA(n_components = 2, kernel="rbf", gamma=0.04)
X_reduced = rbf_pca.fit_transform(X)
```

```

In [29]: from sklearn.decomposition import KernelPCA

lin_pca = KernelPCA(n_components = 2, kernel="linear", fit_inverse_transform=True)
rbf_pca = KernelPCA(n_components = 2, kernel="rbf", gamma=0.0433, fit_inverse_transform=True)
sig_pca = KernelPCA(n_components = 2, kernel="sigmoid", gamma=0.001, coef0=1, fit_inverse_transform=True)

y = t > 6.9


plt.figure(figsize=(11, 4))
for subplot, pca, title in ((131, lin_pca, "Linear kernel"), (132, rbf_pca, "RBF kernel,  $\gamma=0.04$ "), (133, sig_pca, "Sigmoid kernel,  $\gamma=10^{-3}$ ,  $r=1$ ")):
    X_reduced = pca.fit_transform(X)
    if subplot == 132:
        X_reduced_rbf = X_reduced

    plt.subplot(subplot)
    #plt.plot(X_reduced[y, 0], X_reduced[y, 1], "gs")
    #plt.plot(X_reduced[~y, 0], X_reduced[~y, 1], "y^")
    plt.title(title, fontsize=14)
    plt.scatter(X_reduced[:, 0], X_reduced[:, 1], c=t, cmap=plt.cm.hot)
    plt.xlabel(" $z_1$ ", fontsize=18)
    if subplot == 131:
        plt.ylabel(" $z_2$ ", fontsize=18, rotation=0)
    plt.grid(True)

save_fig("kernel_pca_plot")
plt.show()

```

Saving figure kernel_pca_plot




```
In [30]: plt.figure(figsize=(6, 5))

X_inverse = rbf_pca.inverse_transform(X_reduced_rbf)

ax = plt.subplot(111, projection='3d')
ax.view_init(10, -70)
ax.scatter(X_inverse[:, 0], X_inverse[:, 1], X_inverse[:, 2], c=t, cmap=plt.
cm.hot, marker="x")
ax.set_xlabel("")
ax.set_ylabel("")
ax.set_zlabel("")
ax.set_xticklabels([])
ax.set_yticklabels([])
ax.set_zticklabels([])


save_fig("preimage_plot", tight_layout=False)
plt.show()
```

Saving figure preimage_plot



```
In [31]: X_reduced = rbf_pca.fit_transform(X)

plt.figure(figsize=(11, 4))
plt.subplot(132)
plt.scatter(X_reduced[:, 0], X_reduced[:, 1], c=t, cmap=plt.cm.hot, marker="
x")
plt.xlabel("$z_1$", fontsize=18)
plt.ylabel("$z_2$", fontsize=18, rotation=0)
plt.grid(True)
```



```
In [32]: from sklearn.model_selection import GridSearchCV
from sklearn.linear_model import LogisticRegression
from sklearn.pipeline import Pipeline

clf = Pipeline([
    ("kpc", KernelPCA(n_components=2)),
    ("log_reg", LogisticRegression(solver="lbfgs"))
])

param_grid = [{
    "kpc__gamma": np.linspace(0.03, 0.05, 10),
    "kpc__kernel": ["rbf", "sigmoid"]
}]

grid_search = GridSearchCV(clf, param_grid, cv=3)
grid_search.fit(X, y)
```

```
Out[32]: GridSearchCV(cv=3,
                    estimator=Pipeline(steps=[('kpc', KernelPCA(n_components=2)),
                                             ('log_reg', LogisticRegression())]),
                    param_grid=[{'kpc__gamma': array([0.03, 0.03222222, 0.0344444, 0.03666667, 0.03888889, 0.04111111, 0.04333333, 0.04555556, 0.04777778, 0.05]),
                                'kpc__kernel': ['rbf', 'sigmoid']})
```

```
In [33]: print(grid_search.best_params_)

{'kpc__gamma': 0.04333333333333335, 'kpc__kernel': 'rbf'}
```

```
In [34]: rbf_pca = KernelPCA(n_components = 2, kernel="rbf", gamma=0.0433,
                             fit_inverse_transform=True)
X_reduced = rbf_pca.fit_transform(X)
X_preimage = rbf_pca.inverse_transform(X_reduced)
```

```
In [35]: from sklearn.metrics import mean_squared_error

mean_squared_error(X, X_preimage)
```

```
Out[35]: 9.733964708814549e-27
```

LLE

```
In [36]: X, t = make_swiss_roll(n_samples=1000, noise=0.2, random_state=41)
```


```
In [37]: from sklearn.manifold import LocallyLinearEmbedding

lle = LocallyLinearEmbedding(n_components=2, n_neighbors=10, random_state=42)
X_reduced = lle.fit_transform(X)
```

```
In [38]: plt.title("Unrolled swiss roll using LLE", fontsize=14)
plt.scatter(X_reduced[:, 0], X_reduced[:, 1], c=t, cmap=plt.cm.hot)
plt.xlabel("$z_1$", fontsize=18)
plt.ylabel("$z_2$", fontsize=18)
plt.axis([-0.065, 0.055, -0.1, 0.12])
plt.grid(True)

save_fig("lle_unrolling_plot")
plt.show()
```

Saving figure lle_unrolling_plot



MDS, Isomap and t-SNE

```
In [39]: from sklearn.manifold import MDS

mds = MDS(n_components=2, random_state=42)
X_reduced_mds = mds.fit_transform(X)
```

```
In [40]: from sklearn.manifold import Isomap

isomap = Isomap(n_components=2)
X_reduced_isomap = isomap.fit_transform(X)
```

```
In [41]: from sklearn.manifold import TSNE

tsne = TSNE(n_components=2, random_state=42)
X_reduced_tsne = tsne.fit_transform(X)
```

```
In [42]: titles = ["MDS", "Isomap", "t-SNE"]

plt.figure(figsize=(11,4))

for subplot, title, X_reduced in zip((131, 132, 133), titles,
                                     (X_reduced_mds, X_reduced_isomap, X_reduced_tsne)):
    plt.subplot(subplot)
    plt.title(title, fontsize=14)
    plt.scatter(X_reduced[:, 0], X_reduced[:, 1], c=t, cmap=plt.cm.hot)
    plt.xlabel("$z_1$", fontsize=18)
    if subplot == 131:
        plt.ylabel("$z_2$", fontsize=18, rotation=0)
    plt.grid(True)

save_fig("other_dim_reduction_plot")
plt.show()
```

Saving figure other_dim_reduction_plot

