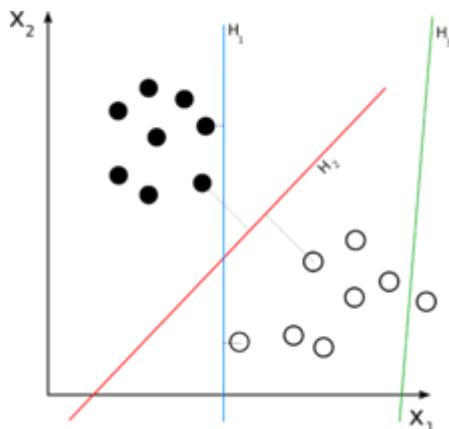


▼ 서포트 벡터 머신(Support Vector Machines)

- 회귀, 분류, 이상치 탐지 등에 사용되는 지도학습 방법
- 클래스 사이의 경계에 위치한 데이터 포인트를 서포트 벡터(support vector)라고 함
- 각 서포트 벡터가 클래스 사이의 결정 경계를 구분하는데 얼마나 중요한지를 학습
- 각 서포트 벡터 사이의 마진이 가장 큰 방향으로 학습
- 서포트 벡터 까지의 거리와 서포트 벡터의 중요도를 기반으로 예측을 수행



- H3은 두 클래스의 점들을 제대로 분류하고 있지 않음
- H1과 H2는 두 클래스의 점들을 분류하는데, H2가 H1보다 더 큰 마진을 갖고 분류하는 것을 확인할 수 있음

```
import multiprocessing
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
plt.style.use(['seaborn-whitegrid'])
```

```
from sklearn.svm import SVC
from sklearn.datasets import load_boston, load_diabetes
from sklearn.datasets import load_breast_cancer, load_iris, load_wine
from sklearn.pipeline import make_pipeline, Pipeline
from sklearn.model_selection import train_test_split, cross_validate, GridSearchCV
from sklearn.preprocessing import StandardScaler, MinMaxScaler
from sklearn.manifold import TSNE
```

▼ SVM을 이용한 회귀 모델과 분류 모델

▼ SVM을 사용한 회귀 모델 (SVR)

```
X, y = load_boston(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, random_state=123)
```

```
model = SVR()
model.fit(X_train, y_train)

print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.2177283706374875
평가 데이터 점수: 0.13544178468518187
```

▼ SVM을 사용한 분류 모델 (SVC)

```
X, y = load_breast_cancer(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, random_state=123)

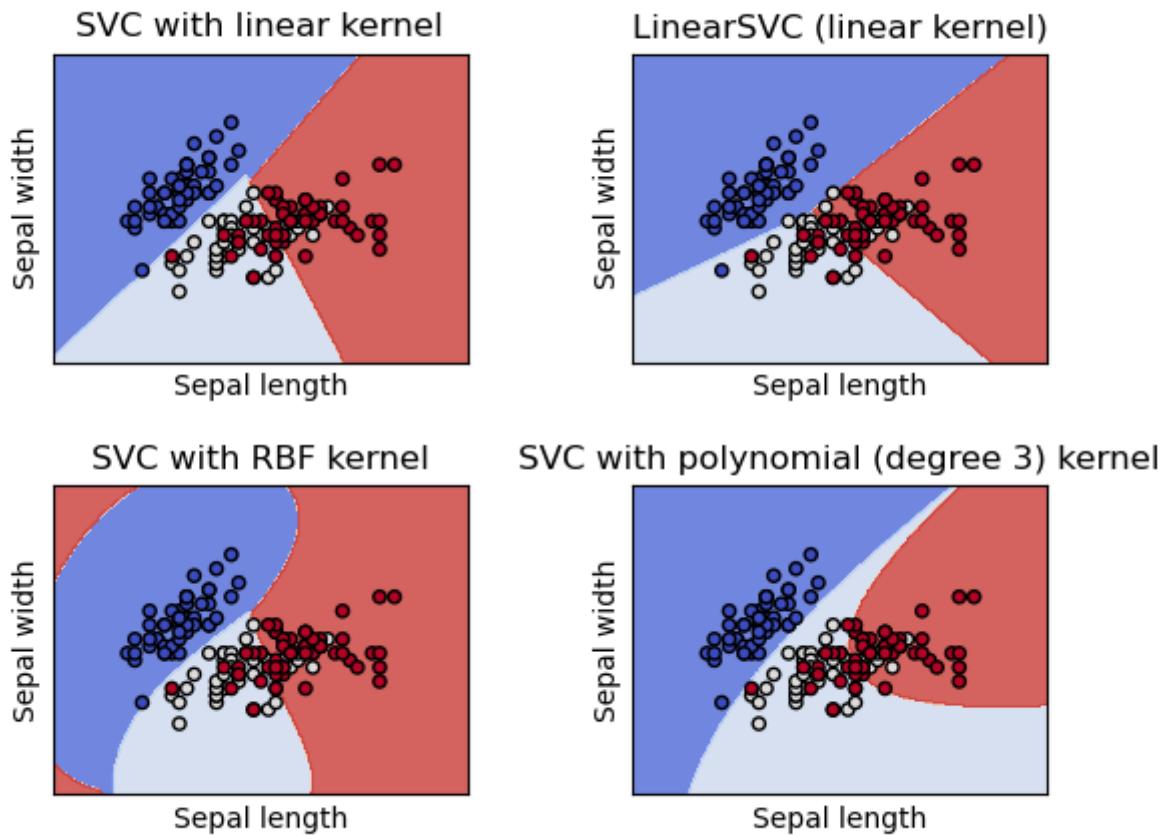
model = SVC()
model.fit(X_train, y_train)

print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.9014084507042254
평가 데이터 점수: 0.9230769230769231
```

▼ 커널 기법

- 입력 데이터를 고차원 공간에 사상해서 비선형 특징을 학습할 수 있도록 확장하는 방법
- scikit-learn에서는 Linear, Polynomial, RBF(Radial Basis Function)등 다양한 커널 기법을 지원



```

X, y = load_boston(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, random_state=123)

linear_svr = SVR(kernel='linear')
linear_svr.fit(X_train, y_train)

print("Linear SVR 학습 데이터 점수: {}".format(linear_svr.score(X_train, y_train)))
print("Linear SVR 평가 데이터 점수: {}".format(linear_svr.score(X_test, y_test)))

polynomial_svr = SVR(kernel='poly')
polynomial_svr.fit(X_train, y_train)

print("Polynomial SVR 학습 데이터 점수: {}".format(polynomial_svr.score(X_train, y_train)))
print("Polynomial SVR 평가 데이터 점수: {}".format(polynomial_svr.score(X_test, y_test)))

rbf_svr = SVR(kernel='rbf')
rbf_svr.fit(X_train, y_train)

print("RBF SVR 학습 데이터 점수: {}".format(rbf_svr.score(X_train, y_train)))
print("RBF SVR 평가 데이터 점수: {}".format(rbf_svr.score(X_test, y_test)))

```

Linear SVR 학습 데이터 점수: 0.715506620496448
 Linear SVR 평가 데이터 점수: 0.6380398541506058
 Polynomial SVR 학습 데이터 점수: 0.2024454261446289
 Polynomial SVR 평가 데이터 점수: 0.133668450367462
 RBF SVR 학습 데이터 점수: 0.2177283706374875
 RBF SVR 평가 데이터 점수: 0.13544178468518187

```

X, y = load_breast_cancer(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, random_state=123)

linear_svc = SVC(kernel='linear')
linear_svc.fit(X_train, y_train)

print("Linear SVC 학습 데이터 점수: {}".format(linear_svc.score(X_train, y_train)))
print("Linear SVC 평가 데이터 점수: {}".format(linear_svc.score(X_test, y_test)))

polynomial_svc = SVC(kernel='poly')
polynomial_svc.fit(X_train, y_train)

print("Polynomial SVC 학습 데이터 점수: {}".format(polynomial_svc.score(X_train, y_train)))
print("Polynomial SVC 평가 데이터 점수: {}".format(polynomial_svc.score(X_test, y_test)))

rbf_svc = SVC(kernel='rbf')
rbf_svc.fit(X_train, y_train)

print("RBF SVC 학습 데이터 점수: {}".format(rbf_svc.score(X_train, y_train)))
print("RBF SVC 평가 데이터 점수: {}".format(rbf_svc.score(X_test, y_test)))

Linear SVC 학습 데이터 점수: 0.960093896713615
Linear SVC 평가 데이터 점수: 0.986013986013986
Polynomial SVC 학습 데이터 점수: 0.9014084507042254
Polynomial SVC 평가 데이터 점수: 0.9230769230769231
RBF SVC 학습 데이터 점수: 0.9014084507042254
RBF SVC 평가 데이터 점수: 0.9230769230769231

```

▼ 매개변수 튜닝

- SVM은 사용하는 커널에 따라 다양한 매개변수 설정 가능
- 매개변수를 변경하면서 성능변화를 관찰

```

X, y = load_breast_cancer(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, random_state=123)

```

```

polynomial_svc = SVC(kernel='poly', degree=2, C=0.1, gamma='auto')
polynomial_svc.fit(X_train, y_train)

print("kernel=poly, degree={}, C={}, gamma={}".format(2, 0.1, 'auto'))
print("Polynomial SVC 학습 데이터 점수: {}".format(polynomial_svc.score(X_train, y_train)))
print("Polynomial SVC 평가 데이터 점수: {}".format(polynomial_svc.score(X_test, y_test)))

kernel=poly, degree=2, C=0.1, gamma=auto
Polynomial SVC 학습 데이터 점수: 0.9835680751173709
Polynomial SVC 평가 데이터 점수: 0.993006993006993

```

```

rbf_svc = SVC(kernel='rbf', C=2.0, gamma='scale')
rbf_svc.fit(X_train, y_train)

print("kernel=poly, C={}, gamma={}".format(2.0, 'scale'))

```

```
print("RBF SVC 학습 데이터 점수: {}".format(rbf_svc.score(X_train, y_train)))
print("RBF SVC 평가 데이터 점수: {}".format(rbf_svc.score(X_test, y_test)))
```

```
kernel=poly, C=2.0, gamma=scale
RBF SVC 학습 데이터 점수: 0.9154929577464789
RBF SVC 평가 데이터 점수: 0.9370629370629371
```

▼ 데이터 전처리

- SVM은 입력 데이터가 정규화되어야 좋은 성능을 보임
- 주로 모든 특성 값을 [0, 1] 범위로 맞추는 방법을 사용
- scikit-learn의 StandardScaler 또는 MinMaxScaler를 사용해 정규화

```
X, y = load_breast_cancer(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, random_state=123)
```

```
model = SVC()
model.fit(X_train, y_train)
```

```
print("SVC 학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("SVC 평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
SVC 학습 데이터 점수: 0.9014084507042254
SVC 평가 데이터 점수: 0.9230769230769231
```

```
scaler = StandardScaler()
X_train = scaler.fit_transform(X_train)
X_test = scaler.transform(X_test)
```

```
model = SVC()
model.fit(X_train, y_train)
```

```
print("SVC 학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("SVC 평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
SVC 학습 데이터 점수: 0.9835680751173709
SVC 평가 데이터 점수: 0.986013986013986
```

```
scaler = MinMaxScaler()
X_train = scaler.fit_transform(X_train)
X_test = scaler.transform(X_test)
```

```
model = SVC()
model.fit(X_train, y_train)
```

```
print("SVC 학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("SVC 평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
SVC 학습 데이터 점수: 0.9812206572769953
```

SVC 평가 데이터 점수: 0.986013986013986

▼ Linear SVR

▼ 보스턴 주택 가격

```
X, y = load_boston(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

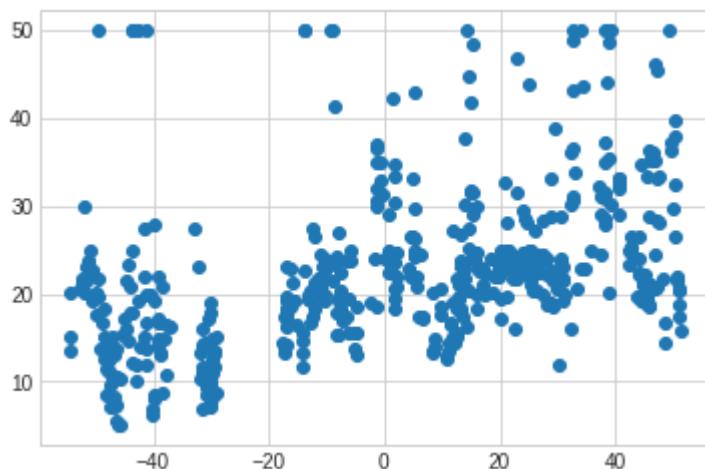
```
model = SVR(kernel='linear')
model.fit(X_train, y_train)
```

```
SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=0.1, gamma='scale',
     kernel='linear', max_iter=-1, shrinking=True, tol=0.001, verbose=False)
```

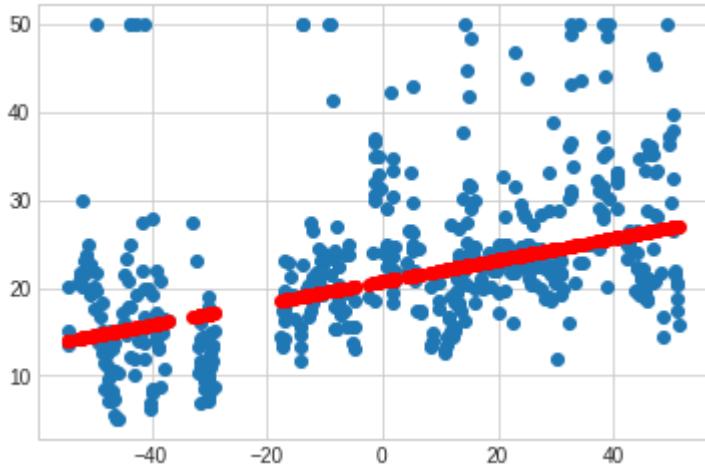
```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

학습 데이터 점수: 0.7426570829547519
평가 데이터 점수: 0.60456758498369

```
X_comp = TSNE(n_components=1).fit_transform(X)
plt.scatter(X_comp, y);
```



```
model.fit(X_comp, y)
predict = model.predict(X_comp)
plt.scatter(X_comp, y)
plt.scatter(X_comp, predict, color='r');
```



```
estimator = make_pipeline(StandardScaler(), SVR(kernel='linear'))
```

```
cross_validate(  
    estimator=estimator,  
    X=X, y=y,  
    cv=5,  
    n_jobs=multiprocessing.cpu_count(),  
    verbose=True  
)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done  5 out of  5 | elapsed:  0.1s finished
{'fit_time': array([0.05558658, 0.02647662, 0.02611208, 0.04099965, 0.02650404]),  

 'score_time': array([0.0020628 , 0.00181913, 0.00179195, 0.00179052, 0.00186872]),  

 'test_score': array([0.76908568, 0.72180141, 0.56428426, 0.14083339, 0.07810211])}
```

```
pipe = Pipeline([('scaler', StandardScaler()),  
                 ('model', SVR(kernel='linear'))])
```

```
param_grid = [{ 'model__gamma': ['scale', 'auto'],
                'model__C': [1.0, 0.1, 0.01],
                'model__epsilon': [1.0, 0.1, 0.01]}]
```

```
gs = GridSearchCV(  
    estimator=pipe,  
    param_grid=param_grid,  
    n_jobs=multiprocessing.cpu_count(),  
    cv=5,  
    verbose=True  
)
```

```
('model',
    SVR(C=1.0, cache_size=200, coef0=0.0,
        degree=3, epsilon=0.1,
        gamma='scale', kernel='linear',
        max_iter=-1, shrinking=True,
        tol=0.001, verbose=False)),
    verbose=False),
iid='deprecated', n_jobs=2,
param_grid=[{'model__C': [1.0, 0.1, 0.01],
    'model__epsilon': [1.0, 0.1, 0.01],
    'model__gamma': ['scale', 'auto']}],
pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
scoring=None, verbose=True)
```

```
gs.best_estimator_
```

```
Pipeline(memory=None,
    steps=[('scaler',
        StandardScaler(copy=True, with_mean=True, with_std=True)),
    ('model',
        SVR(C=0.1, cache_size=200, coef0=0.0, degree=3, epsilon=1.0,
            gamma='scale', kernel='linear', max_iter=-1,
            shrinking=True, tol=0.001, verbose=False)),
    verbose=False)
```

▼ 당뇨병

```
X, y = load_diabetes(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

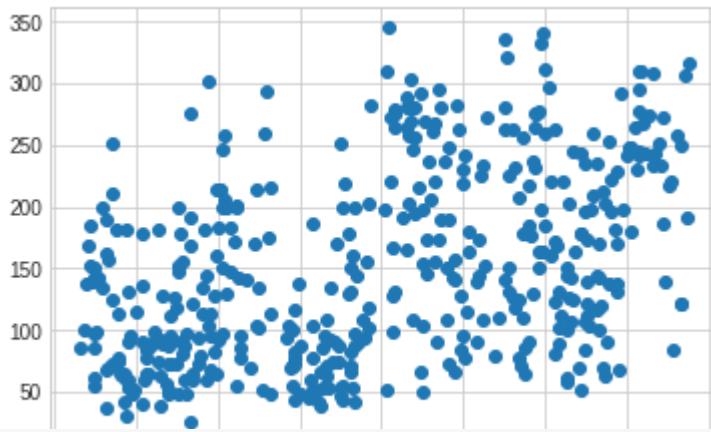
```
model = SVR(kernel='linear')
model.fit(X_train, y_train)
```

```
SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=0.1, gamma='scale',
    kernel='linear', max_iter=-1, shrinking=True, tol=0.001, verbose=False)
```

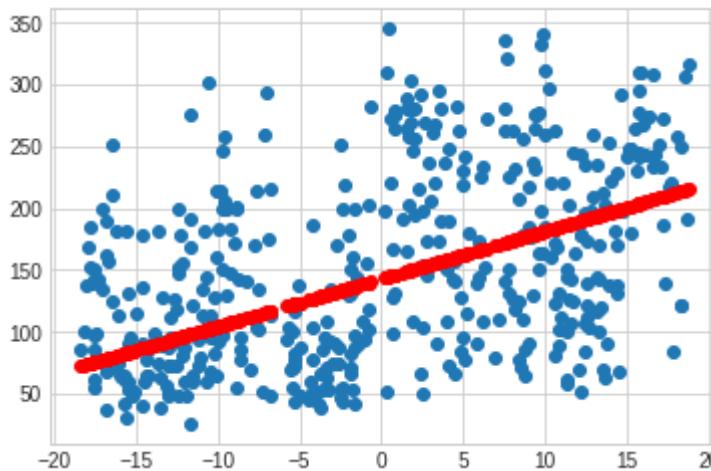
```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.5213639419150948
평가 데이터 점수: 0.4555664384377234
```

```
X_comp = TSNE(n_components=1).fit_transform(X)
plt.scatter(X_comp, y);
```



```
model.fit(X_comp, y)
predict = model.predict(X_comp)
plt.scatter(X_comp, y)
plt.scatter(X_comp, predict, color='r');
```



```
estimator = make_pipeline(StandardScaler(), SVR(kernel='linear'))

cross_validate(
    estimator=estimator,
    X=X, y=y,
    cv=5,
    n_jobs=multiprocessing.cpu_count(),
    verbose=True
)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done   5 out of   5 | elapsed:   0.1s finished
{'fit_time': array([0.0300703, 0.01393008, 0.01182771, 0.02512264, 0.0188756]),
 'score_time': array([0.00177455, 0.00323796, 0.00215983, 0.00144267, 0.00144386]),
 'test_score': array([0.43037242, 0.51653341, 0.48275332, 0.42247202, 0.53076481])}
```

```
pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVR(kernel='linear'))])

param_grid = [{"model__gamma": ['scale', 'auto'],
               "model__C": [1.0, 0.1, 0.01],
               "model__epsilon": [1.0, 0.1, 0.01]}]
```

```
gs = GridSearchCV(  
    estimator=pipe,  
    param_grid=param_grid,  
    n_jobs=multiprocessing.cpu_count(),  
    cv=5,  
    verbose=True  
)
```

```
gs.fit(X, y)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.  
Fitting 5 folds for each of 18 candidates, totalling 90 fits  
[Parallel(n_jobs=2)]: Done 90 out of 90 | elapsed: 0.9s finished  
GridSearchCV(cv=5, error_score='nan',  
    estimator=Pipeline(memory=None,  
        steps=[('scaler',  
            StandardScaler(copy=True,  
                with_mean=True,  
                with_std=True)),  
        ('model',  
            SVR(C=1.0, cache_size=200, coef0=0.0,  
                degree=3, epsilon=0.1,  
                gamma='scale', kernel='linear',  
                max_iter=-1, shrinking=True,  
                tol=0.001, verbose=False))],  
        verbose=False),  
    iid='deprecated', n_jobs=2,  
    param_grid=[{'model__C': [1.0, 0.1, 0.01],  
        'model__epsilon': [1.0, 0.1, 0.01],  
        'model__gamma': ['scale', 'auto']}],  
    pre_dispatch='2*n_jobs', refit=True, return_train_score=False,  
    scoring=None, verbose=True)
```

```
gs.best_estimator_
```

```
Pipeline(memory=None,  
    steps=[('scaler',  
        StandardScaler(copy=True, with_mean=True, with_std=True)),  
        ('model',  
            SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=1.0,  
                gamma='scale', kernel='linear', max_iter=-1,  
                shrinking=True, tol=0.001, verbose=False))],  
    verbose=False)
```

▼ Kernel SVR

▼ 보스턴 주택 가격

```
X, y = load_boston(return_X_y=True)  
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()  
scaler.fit(X_train)
```

```
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

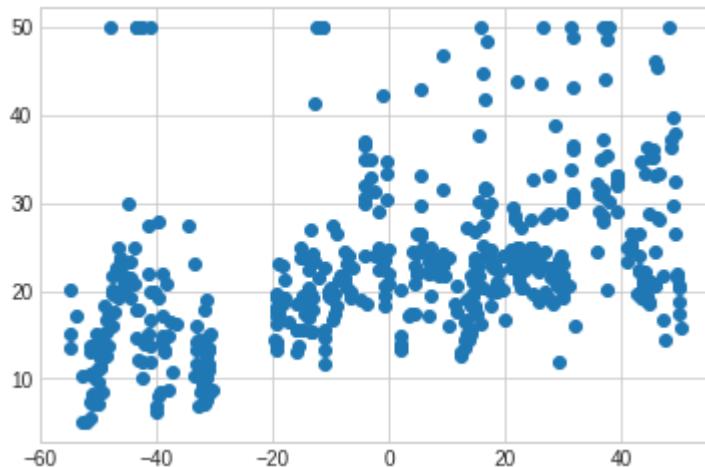
```
model = SVR(kernel='rbf')
model.fit(X_train, y_train)
```

```
SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=0.1, gamma='scale',
     kernel='rbf', max_iter=-1, shrinking=True, tol=0.001, verbose=False)
```

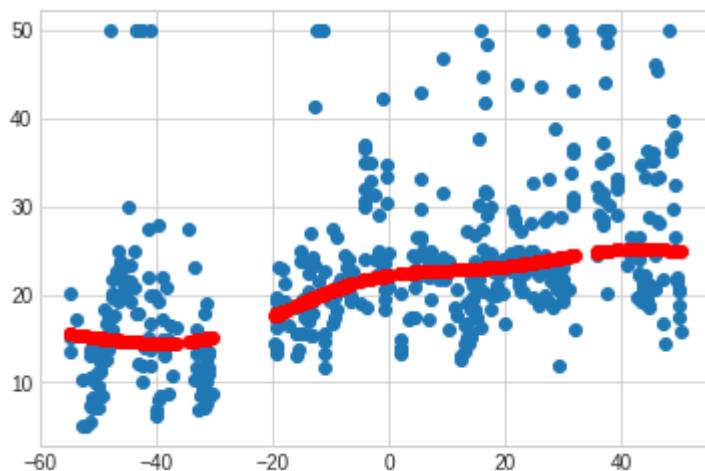
```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.6926959053584525
평가 데이터 점수: 0.5716877813794214
```

```
X_comp = TSNE(n_components=1).fit_transform(X)
plt.scatter(X_comp, y);
```



```
model.fit(X_comp, y)
predict = model.predict(X_comp)
plt.scatter(X_comp, y)
plt.scatter(X_comp, predict, color='r');
```



```
estimator = make_pipeline(StandardScaler(), SVR(kernel='rbf'))
```

```
cross_validate(
    estimator=estimator,
    X=X, y=y,
    cv=5,
    n_jobs=multiprocessing.cpu_count(),
    verbose=True
)
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done  5 out of  5 | elapsed:   0.1s finished
{'fit_time': array([0.04318929, 0.034302 , 0.0437429 , 0.03582788, 0.0250082 ]),
 'score_time': array([0.0056994 , 0.00949693, 0.00315142, 0.00322223, 0.00350928]),
 'test_score': array([ 0.75781445,  0.50211018,  0.04310107,  0.33851703, -0.75997942])}
```

```
pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVR(kernel='rbf'))])
```

```
param_grid = [ {'model__kernel': ['rbf', 'polynomial', 'sigmoid']} ]
```

```
gs = GridSearchCV(
    estimator=pipe,
    param_grid=param_grid,
    n_jobs=multiprocessing.cpu_count(),
    cv=5,
    verbose=True
)
```

```
gs.fit(X, y)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
Fitting 5 folds for each of 3 candidates, totalling 15 fits
[Parallel(n_jobs=2)]: Done  15 out of  15 | elapsed:   0.3s finished
GridSearchCV(cv=5, error_score=nan,
             estimator=Pipeline(memory=None,
                                steps=[('scaler',
                                         StandardScaler(copy=True,
                                                        with_mean=True,
                                                        with_std=True)),
                                       ('model',
                                         SVR(C=1.0, cache_size=200, coef0=0.0,
                                              degree=3, epsilon=0.1,
                                              gamma='scale', kernel='rbf',
                                              max_iter=-1, shrinking=True,
                                              tol=0.001, verbose=False)),
                                       verbose=False),
             iid='deprecated', n_jobs=2,
             param_grid=[ {'model__kernel': ['rbf', 'polynomial', 'sigmoid']} ],
             pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
             scoring=None, verbose=True)
```

```
gs.best_estimator_
```

```
Pipeline(memory=None,
         steps=[('scaler',
                 StandardScaler(copy=True, with_mean=True, with_std=True)),
                ('model',
```

```
        SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=0.1,
             gamma='scale', kernel='rbf', max_iter=-1, shrinking=True,
             tol=0.001, verbose=False)],
            verbose=False)
```

```
pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVR(kernel='rbf'))])

param_grid = [{ 'model__gamma': ['scale', 'auto'],
                'model__C': [1.0, 0.1, 0.01],
                'model__epsilon': [1.0, 0.1, 0.01]}]

gs = GridSearchCV(
    estimator=pipe,
    param_grid=param_grid,
    n_jobs=multiprocessing.cpu_count(),
    cv=5,
    verbose=True
)

gs.fit(X, y)
```

```
Fitting 5 folds for each of 18 candidates, totalling 90 fits
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done  90 out of  90 | elapsed:    1.1s finished
GridSearchCV(cv=5, error_score=nan,
             estimator=Pipeline(memory=None,
                                steps=[('scaler',
                                         StandardScaler(copy=True,
                                                        with_mean=True,
                                                        with_std=True)),
                                       ('model',
                                         SVR(C=1.0, cache_size=200, coef0=0.0,
                                              degree=3, epsilon=0.1,
                                              gamma='scale', kernel='rbf',
                                              max_iter=-1, shrinking=True,
                                              tol=0.001, verbose=False)),
                                         verbose=False),
             iid='deprecated', n_jobs=2,
             param_grid=[{ 'model__C': [1.0, 0.1, 0.01],
                           'model__epsilon': [1.0, 0.1, 0.01],
                           'model__gamma': ['scale', 'auto']}],
             pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
             scoring=None, verbose=True)
```

```
gs.best_estimator_
```

```
Pipeline(memory=None,
          steps=[('scaler',
                  StandardScaler(copy=True, with_mean=True, with_std=True)),
                 ('model',
                   SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=0.01,
                        gamma='scale', kernel='rbf', max_iter=-1, shrinking=True,
                        tol=0.001, verbose=False)),
                 verbose=False)
```

▼ 당뇨병

```
X, y = load_diabetes(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

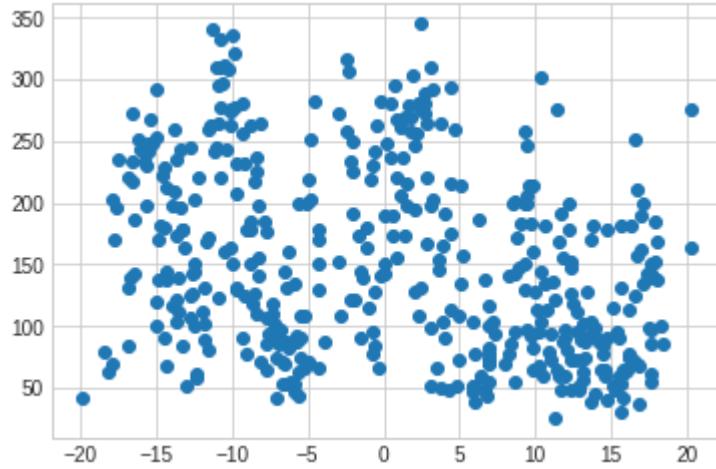
```
model = SVR(kernel='rbf')
model.fit(X_train, y_train)
```

```
SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=0.1, gamma='scale',
     kernel='rbf', max_iter=-1, shrinking=True, tol=0.001, verbose=False)
```

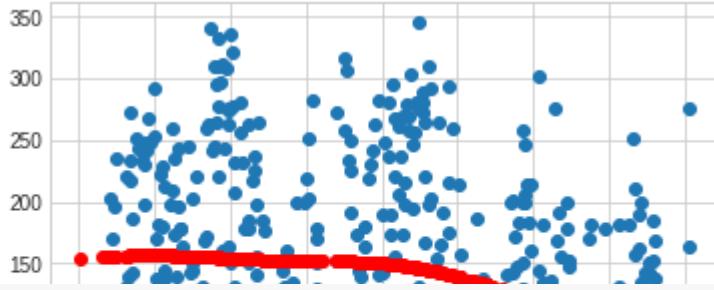
```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.19378724170998662
평가 데이터 점수: 0.13579855194210633
```

```
X_comp = TSNE(n_components=1).fit_transform(X)
plt.scatter(X_comp, y);
```



```
model.fit(X_comp, y)
predict = model.predict(X_comp)
plt.scatter(X_comp, y)
plt.scatter(X_comp, predict, color='r');
```



```
estimator = make_pipeline(StandardScaler(), SVR(kernel='rbf'))
```

```
cross_validate(
    estimator=estimator,
    X=X, y=y,
    cv=5,
    n_jobs=multiprocessing.cpu_count(),
    verbose=True
)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done  5 out of  5 | elapsed:   0.1s finished
{'fit_time': array([0.02605104, 0.01839423, 0.01362157, 0.03073263, 0.01275849]),
 'score_time': array([0.00318933, 0.00282884, 0.00275469, 0.00244069, 0.00260782]),
 'test_score': array([0.14580789, 0.12539919, 0.18163816, 0.12223073, 0.15792085])}
```

```
pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVR(kernel='rbf'))])
```

```
param_grid = [{}{'model_kernel': ['rbf', 'polynomial', 'sigmoid']}]
```

```
gs = GridSearchCV(
    estimator=pipe,
    param_grid=param_grid,
    n_jobs=multiprocessing.cpu_count(),
    cv=5,
    verbose=True
)
```

```
gs.fit(X, y)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
Fitting 5 folds for each of 3 candidates, totalling 15 fits
[Parallel(n_jobs=2)]: Done 12 out of 15 | elapsed:   0.2s remaining:   0.0s
[Parallel(n_jobs=2)]: Done 15 out of 15 | elapsed:   0.2s finished
GridSearchCV(cv=5, error_score=nan,
             estimator=Pipeline(memory=None,
                               steps=[('scaler',
                                       StandardScaler(copy=True,
                                                       with_mean=True,
                                                       with_std=True)),
                                      ('model',
                                       SVR(C=1.0, cache_size=200, coef0=0.0,
                                            degree=3, epsilon=0.1,
                                            gamma='scale', kernel='rbf',
                                            max_iter=-1, shrinking=True,
                                            tol=0.001, verbose=False))],
                               verbose=False),
```

```

    iid='deprecated', n_jobs=2,
param_grid=[{'model__kernel': ['rbf', 'polynomial', 'sigmoid']}],
pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
scoring=None, verbose=True)

gs.best_estimator_


Pipeline(memory=None,
         steps=[('scaler',
                  StandardScaler(copy=True, with_mean=True, with_std=True)),
                ('model',
                  SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=0.1,
                       gamma='scale', kernel='sigmoid', max_iter=-1,
                       shrinking=True, tol=0.001, verbose=False))],
         verbose=False)

pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVR(kernel='sigmoid'))])

param_grid = [{}{'model__gamma': ['scale', 'auto'],
                 'model__C': [1.0, 0.1, 0.01],
                 'model__epsilon': [1.0, 0.1, 0.01]}]

gs = GridSearchCV(
    estimator=pipe,
    param_grid=param_grid,
    n_jobs=multiprocessing.cpu_count(),
    cv=5,
    verbose=True
)
gs.fit(X, y)

Fitting 5 folds for each of 18 candidates, totalling 90 fits
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done 90 out of 90 | elapsed: 1.1s finished
GridSearchCV(cv=5, error_score=nan,
             estimator=Pipeline(memory=None,
                               steps=[('scaler',
                                       StandardScaler(copy=True,
                                                       with_mean=True,
                                                       with_std=True)),
                                      ('model',
                                        SVR(C=1.0, cache_size=200, coef0=0.0,
                                             degree=3, epsilon=0.1,
                                             gamma='scale', kernel='sigmoid',
                                             max_iter=-1, shrinking=True,
                                             tol=0.001, verbose=False))],
                               verbose=False),
             iid='deprecated', n_jobs=2,
             param_grid=[{}{'model__C': [1.0, 0.1, 0.01],
                           'model__epsilon': [1.0, 0.1, 0.01],
                           'model__gamma': ['scale', 'auto']}],
             pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
             scoring=None, verbose=True)

as.best_estimator_

```

```
Pipeline(memory=None,
      steps=[('scaler',
              StandardScaler(copy=True, with_mean=True, with_std=True)),
             ('model',
              SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=1.0,
                   gamma='auto', kernel='sigmoid', max_iter=-1,
                   shrinking=True, tol=0.001, verbose=False))],
      verbose=False)
```

```
model = gs.best_estimator_
model.fit(X_train, y_train)
```

```
Pipeline(memory=None,
      steps=[('scaler',
              StandardScaler(copy=True, with_mean=True, with_std=True)),
             ('model',
              SVR(C=1.0, cache_size=200, coef0=0.0, degree=3, epsilon=1.0,
                   gamma='auto', kernel='sigmoid', max_iter=-1,
                   shrinking=True, tol=0.001, verbose=False))],
      verbose=False)
```

```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.42541451176345024
평가 데이터 점수: 0.27548624665324195
```

▼ Linear SVC

▼ 유방암

```
X, y = load_breast_cancer(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

```
model = SVC(kernel='linear')
model.fit(X_train, y_train)
```

```
SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None, coef0=0.0,
     decision_function_shape='ovr', degree=3, gamma='scale', kernel='linear',
     max_iter=-1, probability=False, random_state=None, shrinking=True,
     tol=0.001, verbose=False)
```

```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
scikit-learn version: 0.23.2 | numpy version: 1.18.5 | python version: 3.7.6
```

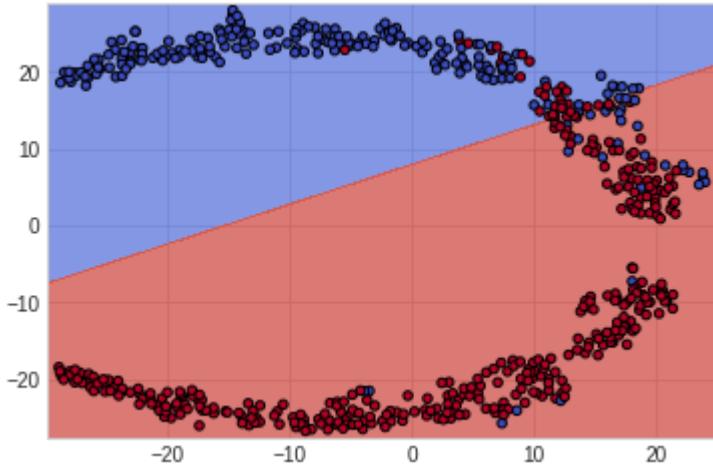
```
학습 데이터 점수: 0.9868131868131869  
평가 데이터 점수: 0.9649122807017544
```

```
def make_meshgrid(x, y, h=.02):  
    x_min, x_max = x.min() - 1, x.max() + 1  
    y_min, y_max = y.min() - 1, y.max() + 1  
    xx, yy = np.meshgrid(np.arange(x_min, x_max, h),  
                         np.arange(y_min, y_max, h))  
  
    return xx, yy
```

```
def plot_contours(clf, xx, yy, **params):  
    Z = clf.predict(np.c_[xx.ravel(), yy.ravel()])  
    Z = Z.reshape(xx.shape)  
    out = plt.contourf(xx, yy, Z, **params)  
  
    return out
```

```
X_comp = TSNE(n_components=2).fit_transform(X)  
X0, X1 = X_comp[:, 0], X_comp[:, 1]  
xx, yy = make_meshgrid(X0, X1)
```

```
model.fit(X_comp, y)  
  
plot_contours(model, xx, yy, cmap=plt.cm.coolwarm, alpha=0.7)  
plt.scatter(X0, X1, c=y, cmap=plt.cm.coolwarm, s=20, edgecolors='k');
```



```
estimator = make_pipeline(StandardScaler(), SVC(kernel='linear'))  
  
cross_validate(  
    estimator=estimator,  
    X=X, y=y,  
    cv=5,  
    n_jobs=multiprocessing.cpu_count(),  
    verbose=True  
)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.  
[Parallel(n_jobs=2)]: Done 5 out of 5 | elapsed: 0.0s finished  
{'fit_time': array([0.00875783, 0.0111444 , 0.00603557, 0.00700355, 0.00548339]),  
 'score_time': array([0.00107265, 0.00106835, 0.00096321, 0.00086117, 0.00064683]),  
 'test_score': array([0.96491228, 0.98245614, 0.96491228, 0.96491228, 0.98230088])}
```

```
pipe = Pipeline([('scaler', StandardScaler()),  
                 ('model', SVC(kernel='linear'))])  
  
param_grid = [{ 'model__gamma': ['scale', 'auto'],  
                'model__C': [1.0, 0.1, 0.01]}]  
  
gs = GridSearchCV(  
    estimator=pipe,  
    param_grid=param_grid,  
    n_jobs=multiprocessing.cpu_count(),  
    cv=5,  
    verbose=True  
)  
  
gs.fit(X, y)
```

```
Fitting 5 folds for each of 6 candidates, totalling 30 fits  
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.  
[Parallel(n_jobs=2)]: Done 30 out of 30 | elapsed: 0.2s finished  
GridSearchCV(cv=5, error_score=nan,  
             estimator=Pipeline(memory=None,  
                               steps=[('scaler',  
                                       StandardScaler(copy=True,  
                                                       with_mean=True,  
                                                       with_std=True)),  
                               ('model',  
                                SVC(C=1.0, break_ties=False,  
                                     cache_size=200, class_weight=None,  
                                     coef0=0.0,  
                                     decision_function_shape='ovr',  
                                     degree=3, gamma='scale',  
                                     kernel='linear', max_iter=-1,  
                                     probability=False,  
                                     random_state=None, shrinking=True,  
                                     tol=0.001, verbose=False)),  
                               verbose=False),  
             iid='deprecated', n_jobs=2,  
             param_grid=[{'model__C': [1.0, 0.1, 0.01],  
                         'model__gamma': ['scale', 'auto']}],  
             pre_dispatch='2*n_jobs', refit=True, return_train_score=False,  
             scoring=None, verbose=True)
```

```
gs.best_estimator_  
  
Pipeline(memory=None,  
         steps=[('scaler',  
                 StandardScaler(copy=True, with_mean=True, with_std=True)),  
                 ('model',  
                  SVC(C=0.1, break_ties=False, cache_size=200, class_weight=None,  
                      coef0=0.0, decision_function_shape='ovr', degree=3,
```

```
gamma='scale', kernel='linear', max_iter=-1,
probability=False, random_state=None, shrinking=True,
tol=0.001, verbose=False)],
verbose=False)
```

▼ 붓꽃

```
X, y = load_iris(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

```
model = SVC(kernel='linear')
model.fit(X_train, y_train)
```

```
SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None, coef0=0.0,
decision_function_shape='ovr', degree=3, gamma='scale', kernel='linear',
max_iter=-1, probability=False, random_state=None, shrinking=True,
tol=0.001, verbose=False)
```

```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.9916666666666667
평가 데이터 점수: 0.9666666666666667
```

```
X_comp = TSNE(n_components=2).fit_transform(X)
X0, X1 = X_comp[:, 0], X_comp[:, 1]
xx, yy = make_meshgrid(X0, X1)
```

```
model.fit(X_comp, y)

plot_contours(model, xx, yy, cmap=plt.cm.coolwarm, alpha=0.7)
plt.scatter(X0, X1, c=y, cmap=plt.cm.coolwarm, s=20, edgecolors='k');
```

```
estimator = make_pipeline(StandardScaler(), SVC(kernel='linear'))\n\ncross_validate(\n    estimator=estimator,\n    X=X, y=y,\n    cv=5,\n    n_jobs=multiprocessing.cpu_count(),\n    verbose=True\n)\n\n[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.\n[Parallel(n_jobs=2)]: Done   5 out of   5 | elapsed:   0.8s finished\n{'fit_time': array([0.00318575, 0.00300145, 0.00215578, 0.00165224, 0.00161457]),\n 'score_time': array([0.00088668, 0.00079918, 0.00066924, 0.00060582, 0.00058913]),\n 'test_score': array([0.96666667, 1.         , 0.93333333, 0.93333333, 1.         ])}\n
```

```
pipe = Pipeline([('scaler', StandardScaler()),  
                 ('model', SVC(kernel='linear'))])
```

```
param_grid = [{model__gamma: ['scale', 'auto'],  
              model__C: [1.0, 0.1, 0.01]}]
```

```
gs = GridSearchCV(  
    estimator=pipe,  
    param_grid=param_grid,  
    n_jobs=multiprocessing.cpu_count(),  
    cv=5,  
    verbose=True  
)
```

```
pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
scoring=None, verbose=True)
```

```
gs.best_estimator_
```

```
Pipeline(memory=None,
         steps=[('scaler',
                  StandardScaler(copy=True, with_mean=True, with_std=True)),
                ('model',
                  SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None,
                       coef0=0.0, decision_function_shape='ovr', degree=3,
                       gamma='scale', kernel='linear', max_iter=-1,
                       probability=False, random_state=None, shrinking=True,
                       tol=0.001, verbose=False))],
         verbose=False)
```

▼ 와인

```
X, y = load_wine(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

```
model = SVC(kernel='linear')
model.fit(X_train, y_train)
```

```
SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None, coef0=0.0,
     decision_function_shape='ovr', degree=3, gamma='scale', kernel='linear',
     max_iter=-1, probability=False, random_state=None, shrinking=True,
     tol=0.001, verbose=False)
```

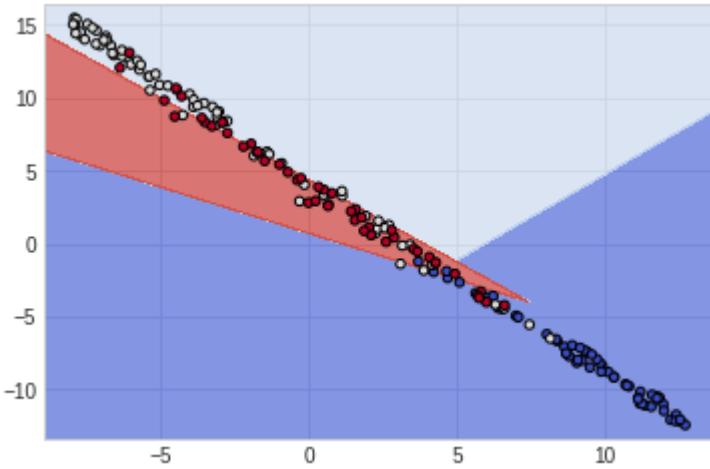
```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 1.0
평가 데이터 점수: 0.9722222222222222
```

```
X_comp = TSNE(n_components=2).fit_transform(X)
X0, X1 = X_comp[:, 0], X_comp[:, 1]
xx, yy = make_meshgrid(X0, X1)
```

```
model.fit(X_comp, y)

plot_contours(model, xx, yy, cmap=plt.cm.coolwarm, alpha=0.7)
plt.scatter(X0, X1, c=y, cmap=plt.cm.coolwarm, s=20, edgecolors='k');
```



```
estimator = make_pipeline(StandardScaler(), SVC(kernel='linear'))
```

```
cross_validate(
    estimator=estimator,
    X=X, y=y,
    cv=5,
    n_jobs=multiprocessing.cpu_count(),
    verbose=True
)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done  5 out of  5 | elapsed:  0.0s finished
{'fit_time': array([0.00331259, 0.00208735, 0.00326419, 0.00264239, 0.00239229]),
 'score_time': array([0.00123405, 0.00052381, 0.00103641, 0.00072551, 0.00060272]),
 'test_score': array([0.94444444, 0.97222222, 0.97222222, 0.97142857, 0.94285714])}
```

```
pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVC(kernel='linear'))])
```

```
param_grid = [ {'model__gamma': ['scale', 'auto'],
                'model__C': [1.0, 0.1, 0.01]} ]
```

```
gs = GridSearchCV(
    estimator=pipe,
    param_grid=param_grid,
    n_jobs=multiprocessing.cpu_count(),
    cv=5,
    verbose=True
)
```

```
gs.fit(X, y)
```

```
Fitting 5 folds for each of 6 candidates, totalling 30 fits
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done  30 out of  30 | elapsed:  0.1s finished
GridSearchCV(cv=5, error_score=nan,
             estimator=Pipeline(memory=None,
                                steps=[('scaler',
                                         StandardScaler(copy=True,
                                                       with_mean=True,
                                                       with_std=True)),
                                ('model',
```

```
SVC(C=1.0, break_ties=False,
    cache_size=200, class_weight=None,
    coef0=0.0,
    decision_function_shape='ovr',
    degree=3, gamma='scale',
    kernel='linear', max_iter=-1,
    probability=False,
    random_state=None, shrinking=True,
    tol=0.001, verbose=False)],
    verbose=False),
iid='deprecated', n_jobs=2,
param_grid=[{'model__C': [1.0, 0.1, 0.01],
             'model__gamma': ['scale', 'auto']},
            pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
            scoring=None, verbose=True)
```

```
gs.best_estimator_
```

```
Pipeline(memory=None,
         steps=[('scaler',
                  StandardScaler(copy=True, with_mean=True, with_std=True)),
                ('model',
                  SVC(C=0.1, break_ties=False, cache_size=200, class_weight=None,
                      coef0=0.0, decision_function_shape='ovr', degree=3,
                      gamma='scale', kernel='linear', max_iter=-1,
                      probability=False, random_state=None, shrinking=True,
                      tol=0.001, verbose=False)),
                verbose=False)
```

▼ Kernel SVC

▼ 유방암

```
X, y = load_breast_cancer(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

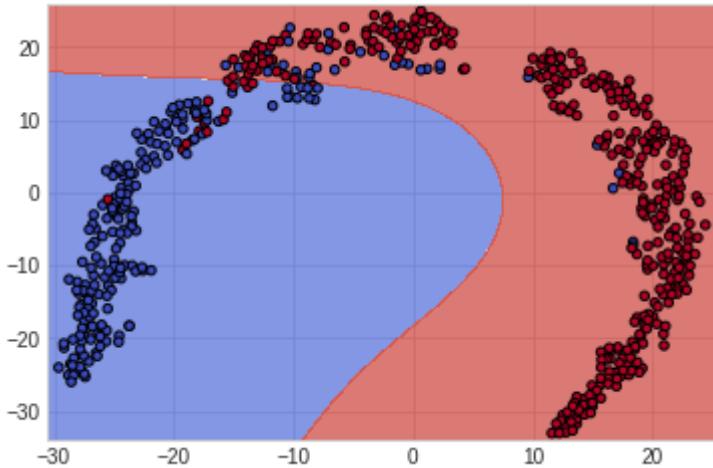
```
model = SVC(kernel='rbf')
model.fit(X_train, y_train)
```

```
SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None, coef0=0.0,
    decision_function_shape='ovr', degree=3, gamma='scale', kernel='rbf',
    max_iter=-1, probability=False, random_state=None, shrinking=True,
    tol=0.001, verbose=False)
```

```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.989010989010989  
평가 데이터 점수: 0.9649122807017544
```

```
X_comp = TSNE(n_components=2).fit_transform(X)  
X0, X1 = X_comp[:, 0], X_comp[:, 1]  
xx, yy = make_meshgrid(X0, X1)  
  
model.fit(X_comp, y)  
  
plot_contours(model, xx, yy, cmap=plt.cm.coolwarm, alpha=0.7)  
plt.scatter(X0, X1, c=y, cmap=plt.cm.coolwarm, s=20, edgecolors='k');
```



```
estimator = make_pipeline(StandardScaler(), SVC(kernel='linear'))  
  
cross_validate(  
    estimator=estimator,  
    X=X, y=y,  
    cv=5,  
    n_jobs=multiprocessing.cpu_count(),  
    verbose=True  
)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.  
[Parallel(n_jobs=2)]: Done 5 out of 5 | elapsed: 0.9s finished  
{'fit_time': array([0.00946522, 0.0070467 , 0.00659204, 0.00564122, 0.00614333]),  
 'score_time': array([0.0012424 , 0.00125408, 0.00098777, 0.00096607, 0.00108624]),  
 'test_score': array([0.96491228, 0.98245614, 0.96491228, 0.96491228, 0.98230088])}
```

```
pipe = Pipeline([('scaler', StandardScaler()),  
                 ('model', SVC(kernel='linear'))])  
  
param_grid = [{ 'model__gamma': ['scale', 'auto'],  
                'model__C': [1.0, 0.1, 0.01]}]  
  
gs = GridSearchCV(  
    estimator=pipe,  
    param_grid=param_grid,  
    n_jobs=multiprocessing.cpu_count(),  
    ...)
```

```

        cv=5,
        verbose=True
    )

gs.fit(X, y)

Fitting 5 folds for each of 6 candidates, totalling 30 fits
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done 30 out of 30 | elapsed: 0.2s finished
GridSearchCV(cv=5, error_score='nan',
             estimator=Pipeline(memory=None,
                                steps=[('scaler',
                                         StandardScaler(copy=True,
                                                        with_mean=True,
                                                        with_std=True)),
                               ('model',
                                SVC(C=1.0, break_ties=False,
                                     cache_size=200, class_weight=None,
                                     coef0=0.0,
                                     decision_function_shape='ovr',
                                     degree=3, gamma='scale',
                                     kernel='linear', max_iter=-1,
                                     probability=False,
                                     random_state=None, shrinking=True,
                                     tol=0.001, verbose=False))],
             iid='deprecated', n_jobs=2,
             param_grid=[{'model__C': [1.0, 0.1, 0.01],
                          'model__gamma': ['scale', 'auto']},
                         pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
                         scoring=None, verbose=True)

```

```
gs.best_estimator_
```

```

Pipeline(memory=None,
         steps=[('scaler',
                 StandardScaler(copy=True, with_mean=True, with_std=True)),
                ('model',
                 SVC(C=0.1, break_ties=False, cache_size=200, class_weight=None,
                      coef0=0.0, decision_function_shape='ovr', degree=3,
                      gamma='scale', kernel='linear', max_iter=-1,
                      probability=False, random_state=None, shrinking=True,
                      tol=0.001, verbose=False))],
         verbose=False)

```

▼ 붓꽃

```

X, y = load_iris(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)

```

```

scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)

```

```
model = SVC(kernel='rbf')
model.fit(X_train, y_train)
```

```
SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None, coef0=0.0,
    decision_function_shape='ovr', degree=3, gamma='scale', kernel='rbf',
    max_iter=-1, probability=False, random_state=None, shrinking=True,
    tol=0.001, verbose=False)
```

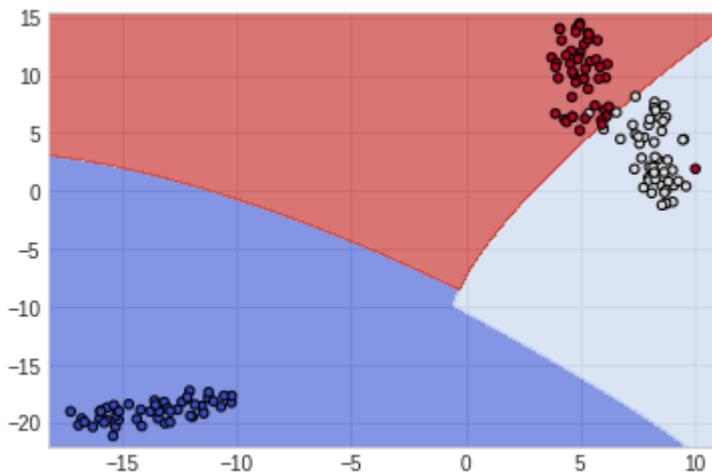
```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 1.0
평가 데이터 점수: 0.9
```

```
X_comp = TSNE(n_components=2).fit_transform(X)
X0, X1 = X_comp[:, 0], X_comp[:, 1]
xx, yy = make_meshgrid(X0, X1)
```

```
model.fit(X_comp, y)
```

```
plot_contours(model, xx, yy, cmap=plt.cm.coolwarm, alpha=0.7)
plt.scatter(X0, X1, c=y, cmap=plt.cm.coolwarm, s=20, edgecolors='k');
```



```
estimator = make_pipeline(StandardScaler(), SVC(kernel='linear'))

cross_validate(
    estimator=estimator,
    X=X, y=y,
    cv=5,
    n_jobs=multiprocessing.cpu_count(),
    verbose=True
)
```

```
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done   5 out of   5 | elapsed:   0.9s finished
{'fit_time': array([0.00280333, 0.00286293, 0.00202656, 0.00177312, 0.00177383]),
 'score_time': array([0.00074673, 0.00067592, 0.0006578 , 0.0006671 , 0.00061393]),
 'test_score': array([0.96666667, 1.          , 0.93333333, 0.93333333, 1.          ])}
```

```

pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVC(kernel='linear'))])

param_grid = [{ 'model__gamma': ['scale', 'auto'],
                'model__C': [1.0, 0.1, 0.01]}]

gs = GridSearchCV(
    estimator=pipe,
    param_grid=param_grid,
    n_jobs=multiprocessing.cpu_count(),
    cv=5,
    verbose=True
)

gs.fit(X, y)

Fitting 5 folds for each of 6 candidates, totalling 30 fits
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done 30 out of 30 | elapsed: 0.2s finished
GridSearchCV(cv=5, error_score='nan',
             estimator=Pipeline(memory=None,
                                steps=[('scaler',
                                         StandardScaler(copy=True,
                                                        with_mean=True,
                                                        with_std=True)),
                                       ('model',
                                         SVC(C=1.0, break_ties=False,
                                              cache_size=200, class_weight=None,
                                              coef0=0.0,
                                              decision_function_shape='ovr',
                                              degree=3, gamma='scale',
                                              kernel='linear', max_iter=-1,
                                              probability=False,
                                              random_state=None, shrinking=True,
                                              tol=0.001, verbose=False)),
                                         verbose=False),
             iid='deprecated', n_jobs=2,
             param_grid=[{'model__C': [1.0, 0.1, 0.01],
                          'model__gamma': ['scale', 'auto']}],
             pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
             scoring=None, verbose=True)

```

```
gs.best_estimator_
```

```

Pipeline(memory=None,
         steps=[('scaler',
                 StandardScaler(copy=True, with_mean=True, with_std=True)),
                ('model',
                  SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None,
                      coef0=0.0, decision_function_shape='ovr', degree=3,
                      gamma='scale', kernel='linear', max_iter=-1,
                      probability=False, random_state=None, shrinking=True,
                      tol=0.001, verbose=False))],
         verbose=False)

```

```
X, y = load_wine(return_X_y=True)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
```

```
scaler = StandardScaler()
scaler.fit(X_train)
X_train = scaler.transform(X_train)
X_test = scaler.transform(X_test)
```

```
model = SVC(kernel='rbf')
model.fit(X_train, y_train)
```

```
SVC(C=1.0, break_ties=False, cache_size=200, class_weight=None, coef0=0.0,
     decision_function_shape='ovr', degree=3, gamma='scale', kernel='rbf',
     max_iter=-1, probability=False, random_state=None, shrinking=True,
     tol=0.001, verbose=False)
```

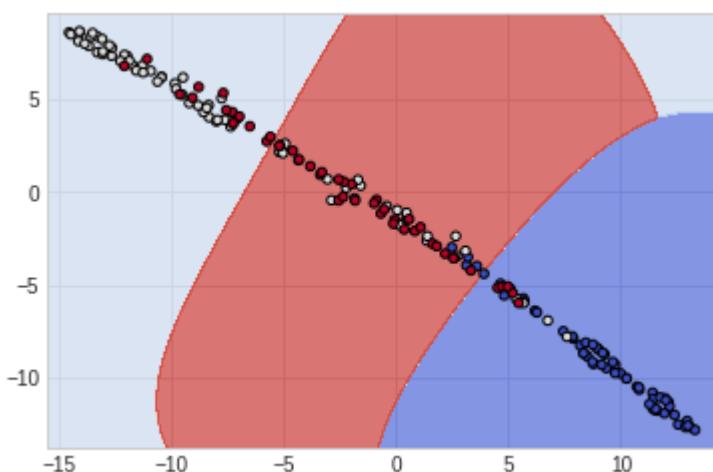
```
print("학습 데이터 점수: {}".format(model.score(X_train, y_train)))
print("평가 데이터 점수: {}".format(model.score(X_test, y_test)))
```

```
학습 데이터 점수: 0.9929577464788732
평가 데이터 점수: 1.0
```

```
X_comp = TSNE(n_components=2).fit_transform(X)
X0, X1 = X_comp[:, 0], X_comp[:, 1]
xx, yy = make_meshgrid(X0, X1)
```

```
model.fit(X_comp, y)

plot_contours(model, xx, yy, cmap=plt.cm.coolwarm, alpha=0.7)
plt.scatter(X0, X1, c=y, cmap=plt.cm.coolwarm, s=20, edgecolors='k');
```



```
estimator = make_pipeline(StandardScaler(), SVC(kernel='linear'))
```

```
cross_validate(
    estimator=estimator,
    X=X, y=y,
    cv=5.
```

```

...,
n_jobs=multiprocessing.cpu_count(),
verbose=True
)

[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done  5 out of  5 | elapsed:  0.9s finished
{'fit_time': array([0.00311828, 0.00332761, 0.00241804, 0.00239849, 0.00272346]),
 'score_time': array([0.00077629, 0.00071621, 0.00069547, 0.00066233, 0.00048327]),
 'test_score': array([0.94444444, 0.97222222, 0.97222222, 0.97142857, 0.94285714])}

pipe = Pipeline([('scaler', StandardScaler()),
                 ('model', SVC(kernel='linear'))])

param_grid = [{('model__gamma': ['scale', 'auto'],
               'model__C': [1.0, 0.1, 0.01])}]

gs = GridSearchCV(
    estimator=pipe,
    param_grid=param_grid,
    n_jobs=multiprocessing.cpu_count(),
    cv=5,
    verbose=True
)

gs.fit(X, y)

Fitting 5 folds for each of 6 candidates, totalling 30 fits
[Parallel(n_jobs=2)]: Using backend LokyBackend with 2 concurrent workers.
[Parallel(n_jobs=2)]: Done 30 out of 30 | elapsed:  0.1s finished
GridSearchCV(cv=5, error_score=nan,
             estimator=Pipeline(memory=None,
                                steps=[('scaler',
                                         StandardScaler(copy=True,
                                                        with_mean=True,
                                                        with_std=True)),
                                       ('model',
                                         SVC(C=1.0, break_ties=False,
                                              cache_size=200, class_weight=None,
                                              coef0=0.0,
                                              decision_function_shape='ovr',
                                              degree=3, gamma='scale',
                                              kernel='linear', max_iter=-1,
                                              probability=False,
                                              random_state=None, shrinking=True,
                                              tol=0.001, verbose=False)),
                                       verbose=False),
             iid='deprecated', n_jobs=2,
             param_grid=[{('model__C': [1.0, 0.1, 0.01],
                           'model__gamma': ['scale', 'auto'])}],
             pre_dispatch='2*n_jobs', refit=True, return_train_score=False,
             scoring=None, verbose=True)

gs.best_estimator_

Pipeline(memory=None,
         steps=[('scaler',
                 StandardScaler(copy=True, with_mean=True, with_std=True)),
```

```
('model',
SVC(C=0.1, break_ties=False, cache_size=200, class_weight=None,
    coef0=0.0, decision_function_shape='ovr', degree=3,
    gamma='scale', kernel='linear', max_iter=-1,
    probability=False, random_state=None, shrinking=True,
    tol=0.001, verbose=False)],
verbose=False)
```

