

**ENTERPRISE ARTIFICIAL
INTELLIGENCE
TRANSFORMATION**

BY RASHED HAQ

Contents

Figure 2.1	4
Figure 2.2	4
Figure 2.3	4
Figure 2.4	5
Figure 2.5	5
Figure 2.6	6
Figure 5.1	7
Figure 5.2	8
Figure 5.3	9
Figure 7.1	10
Figure 8.1	11
Figure 8.2	12
Figure 8.3	13
Figure 8.4	14
Figure 8.5	15
Figure 8.6	16
Figure 8.7	16
Figure 8.8	17
Figure 8.9	18
Figure 8.10	19
Figure 9.1	20
Figure 9.2	21
Figure 9.3	21
Figure 9.4	22
Figure 9.5	22
Figure 10.1	23
Figure 10.2	24
Figure 11.1	25
Figure 11.2	25
Figure 12.1	26
Figure 12.2	27
Figure 12.3	28
Figure 12.4	29
Figure 12.5	29
Figure 12.6	30
Figure 12.7	31

Figure 13.1	32
Figure 13.2	32
Figure 13.3	33
Figure 13.4	33
Figure 13.5	34
Figure 13.6	34
Figure 13.7	34
Figure 13.8	35
Figure 13.9	36
Figure 13.10	37
Figure 13.11	38
Figure 13.12	39
Figure 13.13	40
Figure 13.14	41
Figure 13.15	41
Figure 13.16	41
Figure 13.17	42
Figure 13.18	42
Figure 13.19	42
Figure 13.20	43
Figure 13.21	44
Figure 13.22	45
Figure 13.23	45
Figure 13.24	46

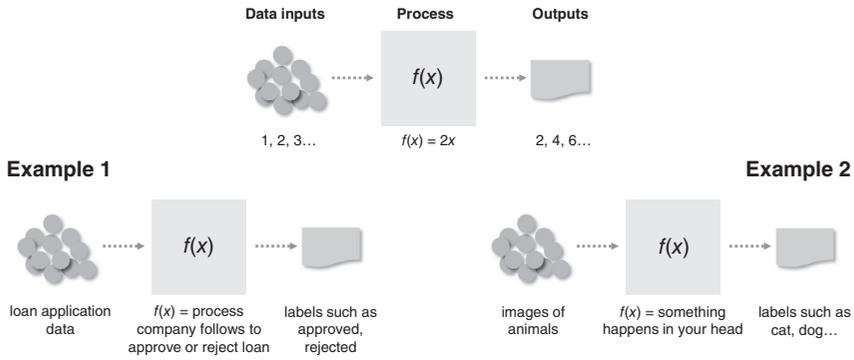


Figure 2.1 Examples of functions $f(x)$ that can be estimated by using machine learning on the input and output datasets.

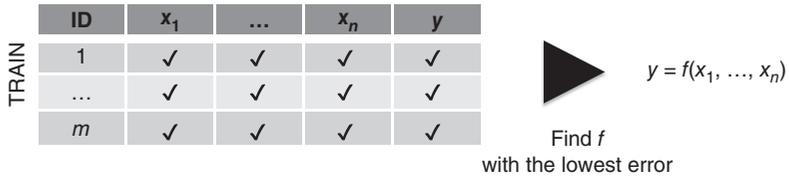


Figure 2.2 Using training data for customers 1 to m to estimate f that will predict y given x_1, \dots, x_n .



Figure 2.3 Using the machine-learning model (f) to predict if customer number $m + 1$ will churn.

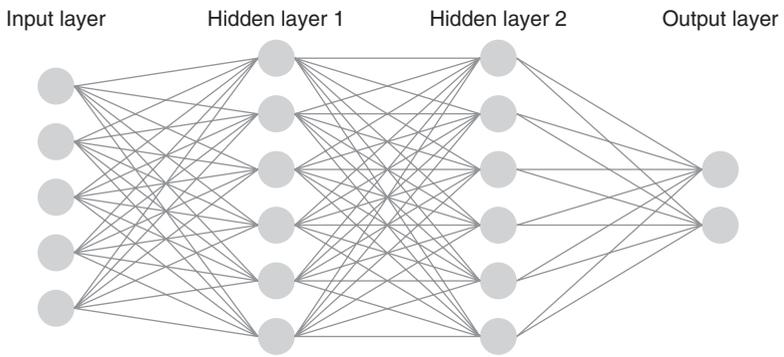


Figure 2.4 An example of a deep neural network.

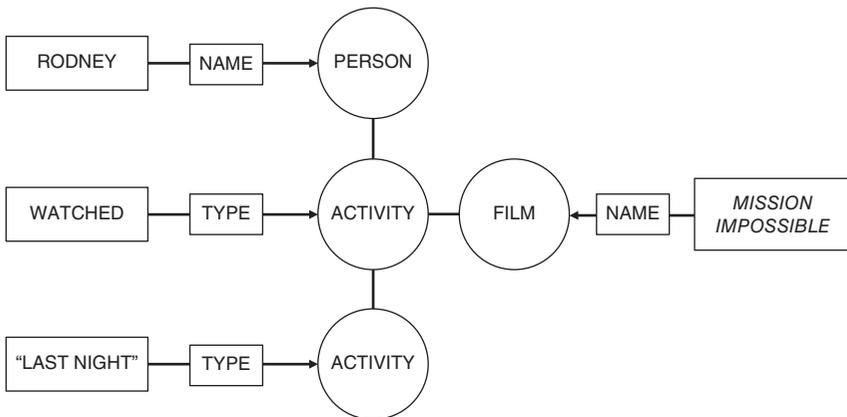


Figure 2.5 Example of a type of knowledge graph.

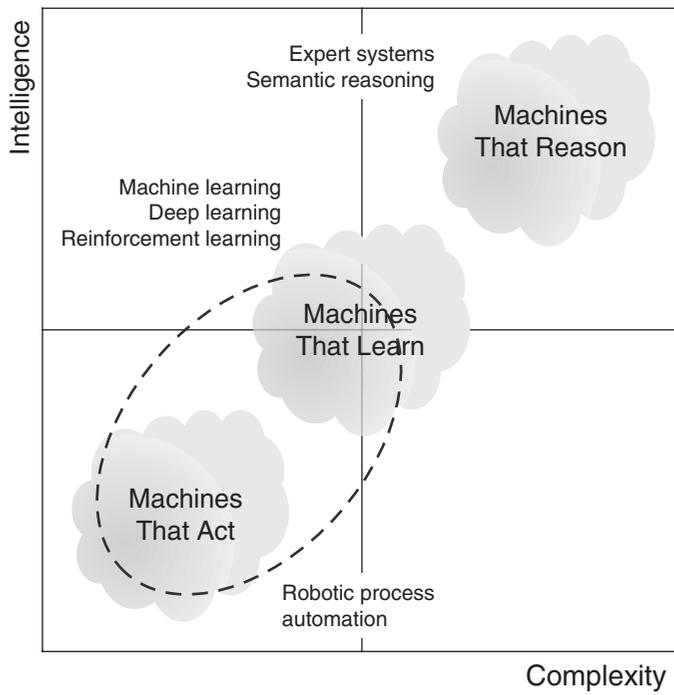


Figure 2.6 Types of AI systems.

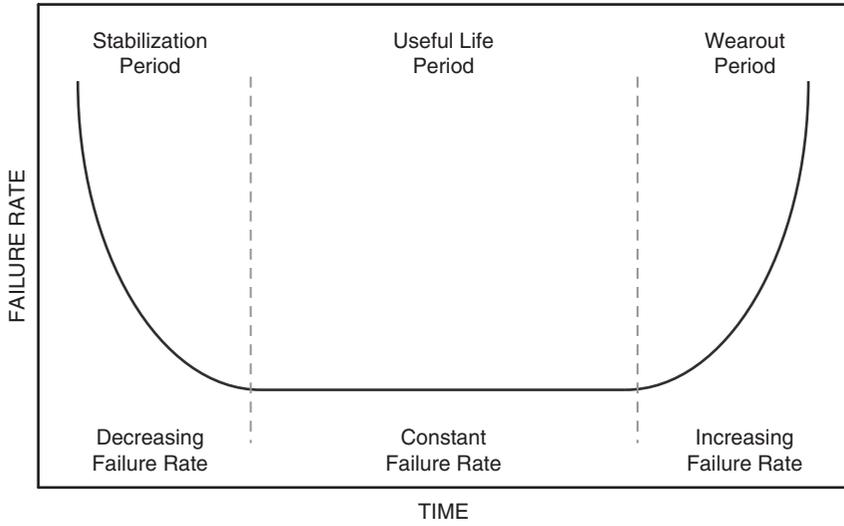


Figure 5.1 Heuristic showing different failure rates during equipment component lifecycle.

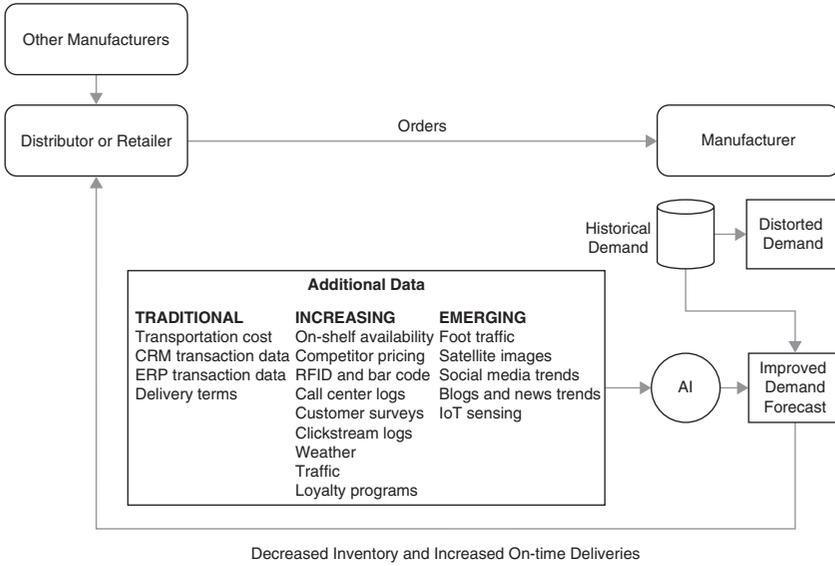


Figure 5.2 Demand forecasting using historical sales and new data sources.



Figure 5.3 Energy trading scenario.

Firmographic	Demographic	Technographic	Relationship	Satellite	News and Blogs	Device	Behavioral
Revenues Employees SIC/NAICS codes Office locations Department and business unit Press releases Job postings	Biography Title/role Time in title/role Office Location Department Education	IT infrastructure Applications (used or installed) Professional services contracts Equipment Other technologies Contract renewal date	Social network Reporting relationship External social contacts Internal colleague and employee contacts Influence rating (Klout or proprietary score)	Foot traffic Parking lot usage Inventory and movement	News sites Analyst and broker reports Blogs	Reverse IP address Geolocation Usage	Content consumption Search history Site, page, and app visits Social posts

Figure 7.1 Different types of third-party data that is available commercially or through the web.

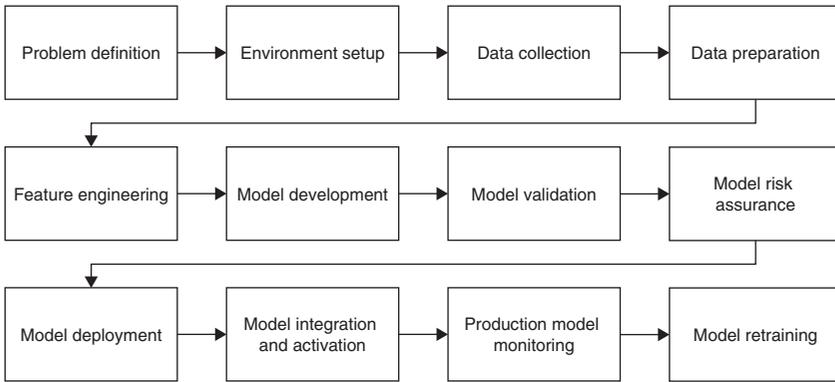


Figure 8.1 The workflow for AI, machine learning, and data science projects.

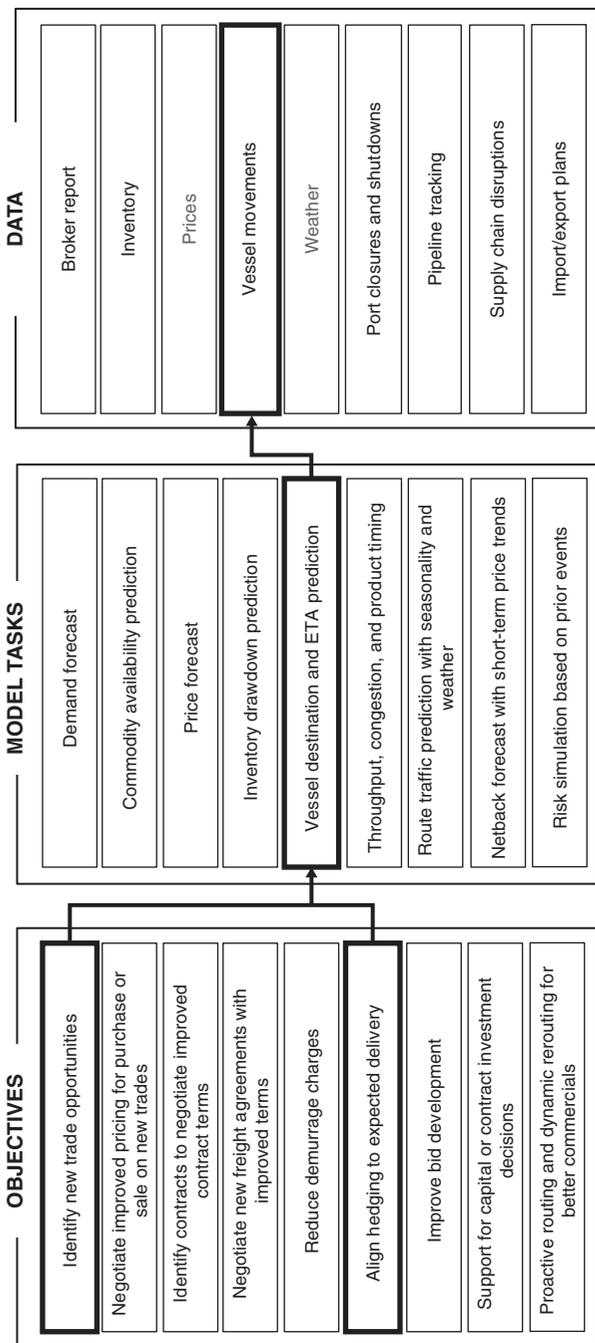


Figure 8.2 A sample map of use case objectives, modeling tasks to support the decisions, and data required to support modeling.

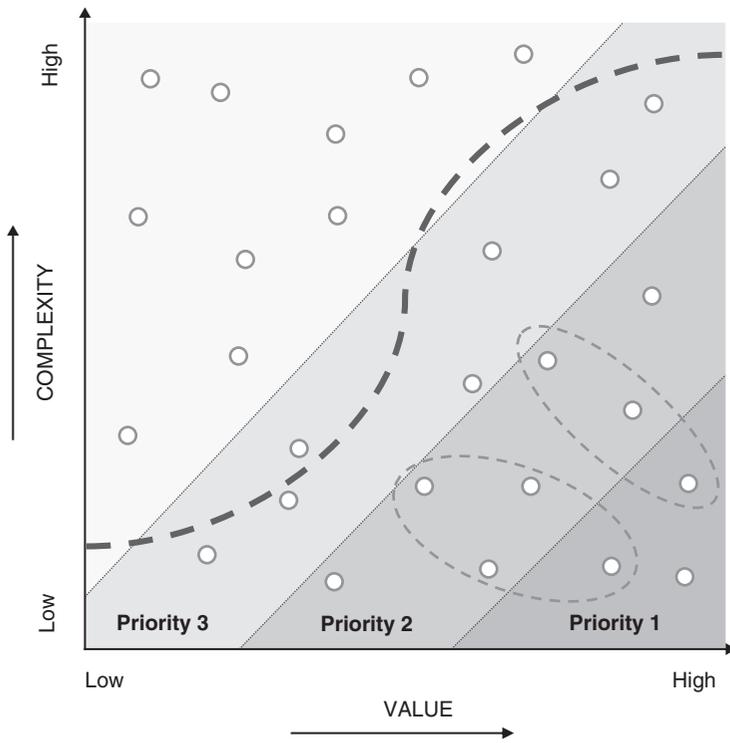


Figure 8.3 Graph showing use cases by value and complexity.

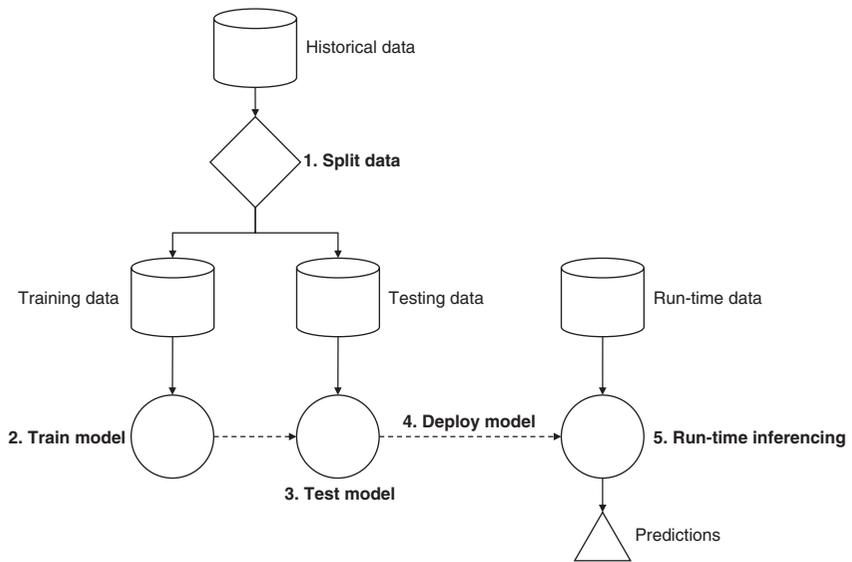


Figure 8.4 Process for training and validating the model.

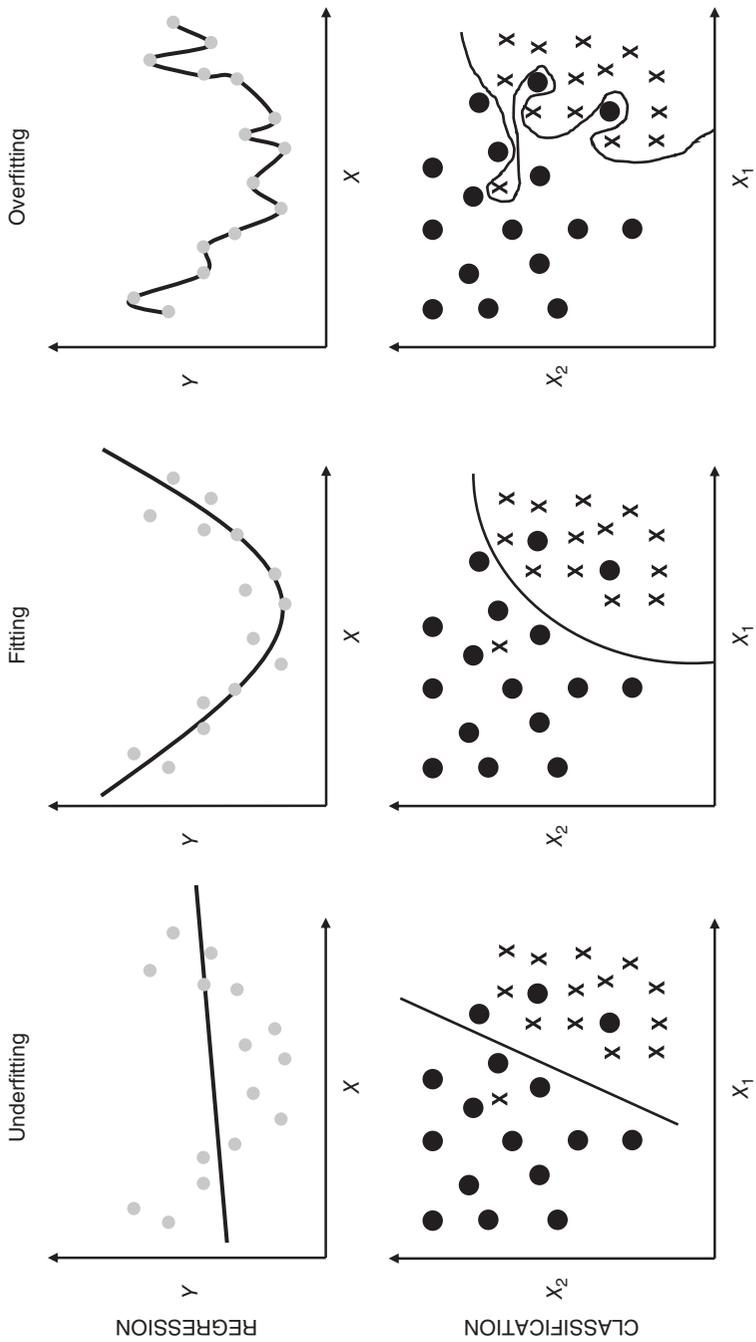


Figure 8.5 Underfitting and overfitting for regression models (top) and for classification models (bottom).

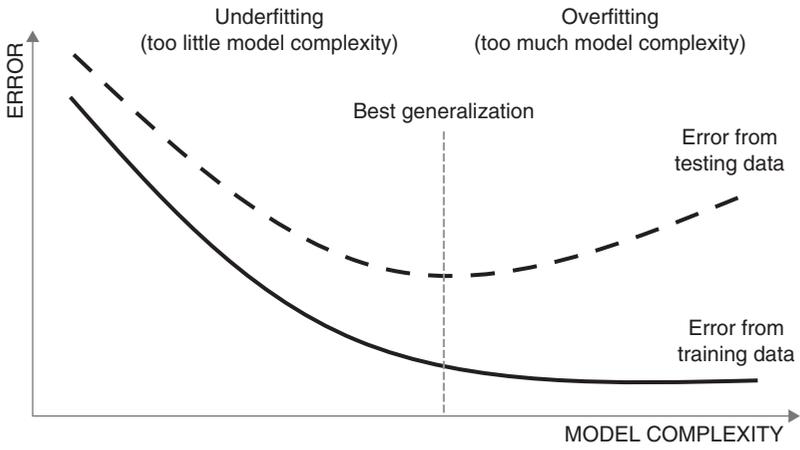


Figure 8.6 Training error versus testing error.

		ACTUAL VALUE	
		Positive	Negative
PREDICTED VALUE	Positive	True Positive	False Positive
	Negative	False Negative	True Negative

Figure 8.7 The confusion matrix setup.

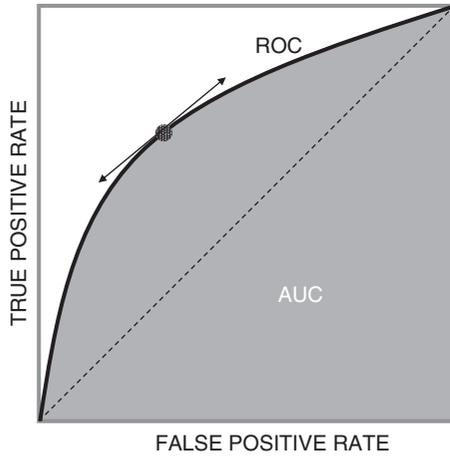


Figure 8.8 Receiver operating characteristics (ROC) curve and the area under the curve (AUC).

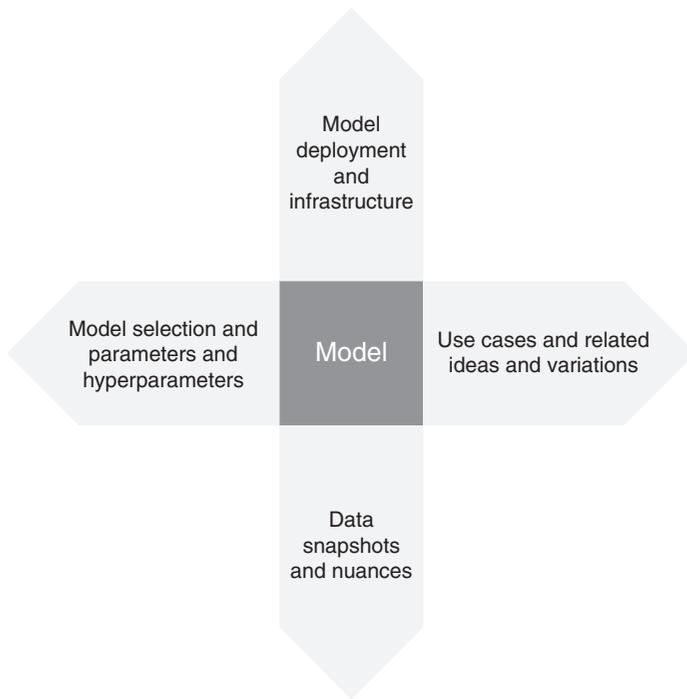


Figure 8.9 Comprehensive model management spans four types of configurations.

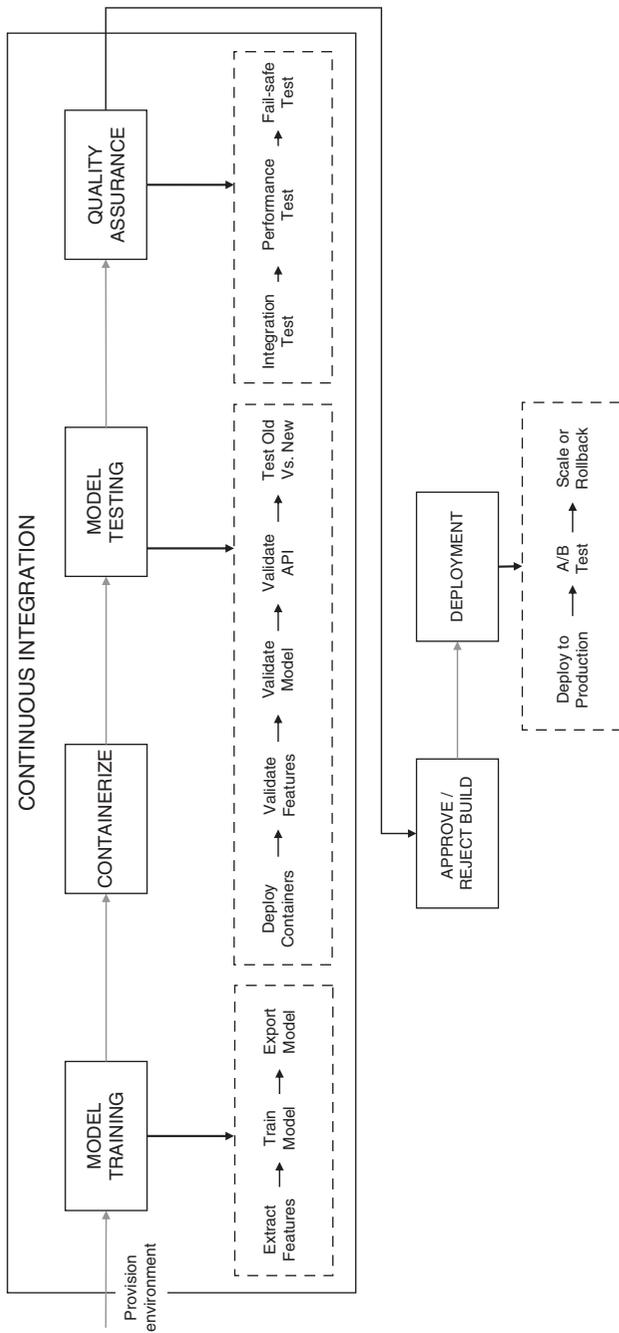


Figure 8.10 AI DevOps process.

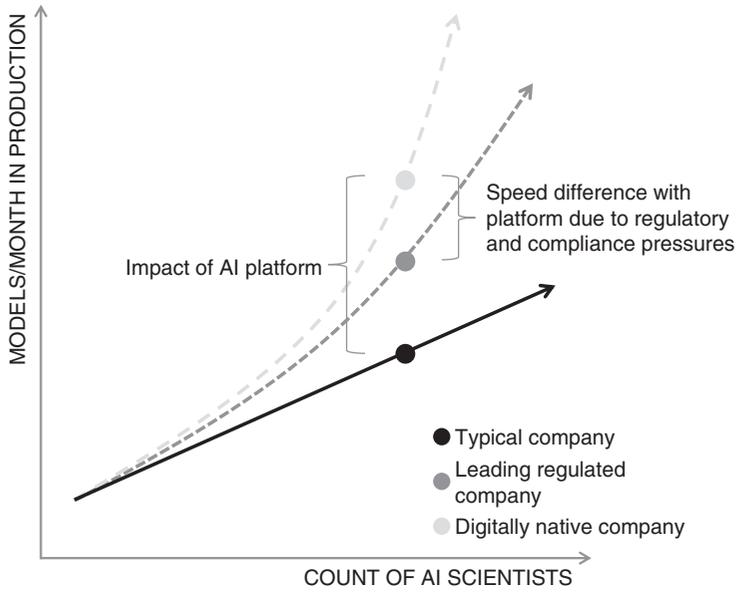


Figure 9.1 Impact of using an AI platform.

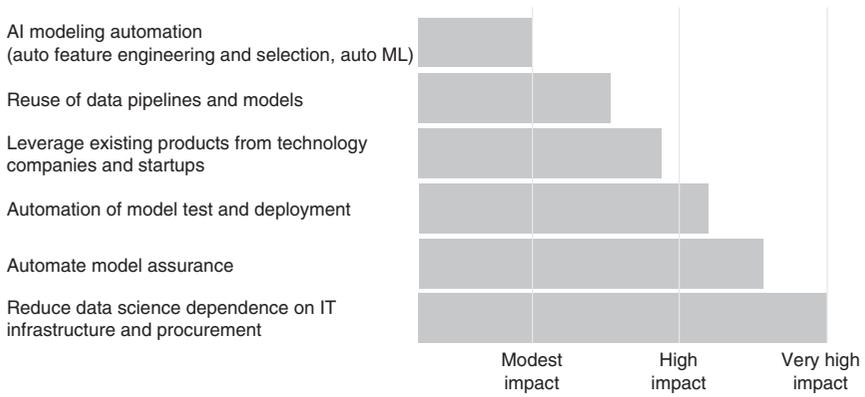


Figure 9.2 Summary of benefits of using an AI platform.

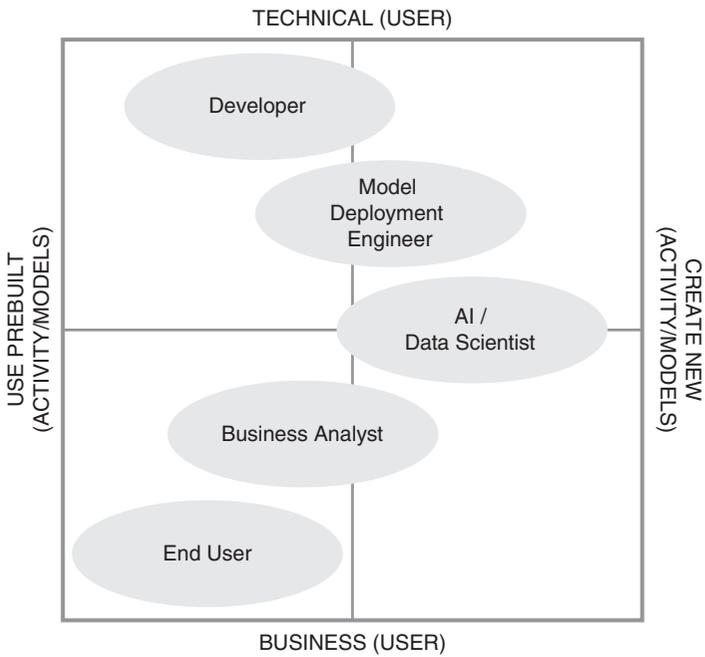


Figure 9.3 Types of users of an AI platform (vertical axis) and how they engage with AI models (horizontal axis).

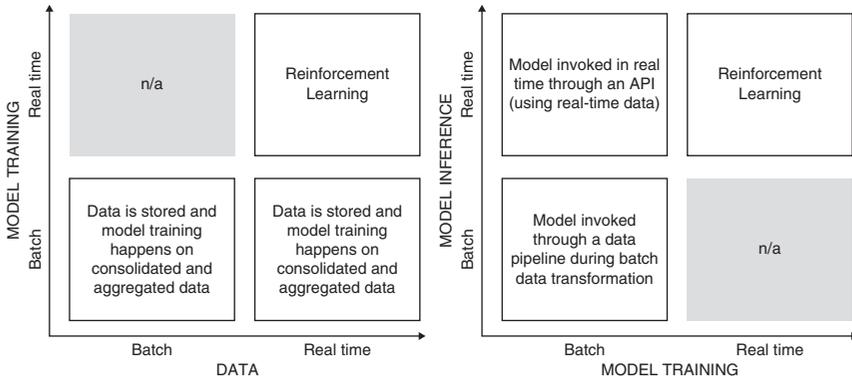


Figure 9.4 Batch versus real time for data, model training, and model inferencing.

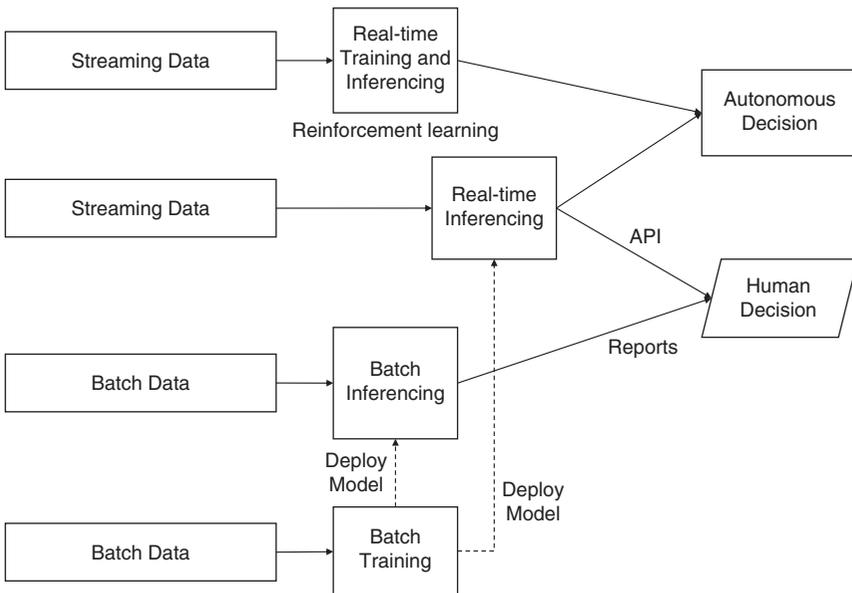


Figure 9.5 The different patterns of batch or streaming data, model training, model inference, and usage.

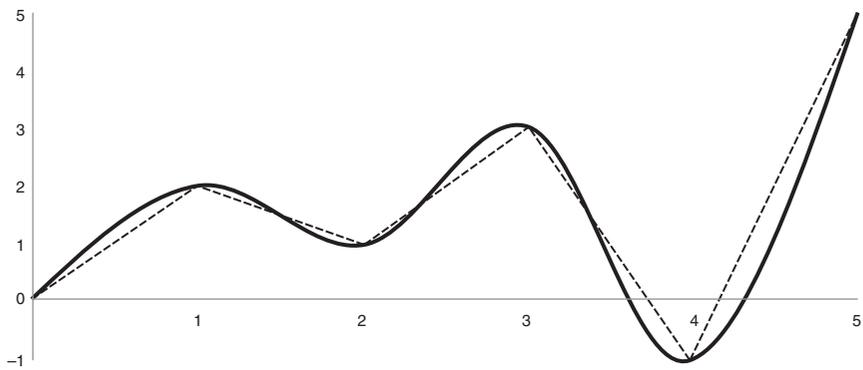


Figure 10.1 Approximating a polynomial function using simpler linear functions in different parts of the x -axis.

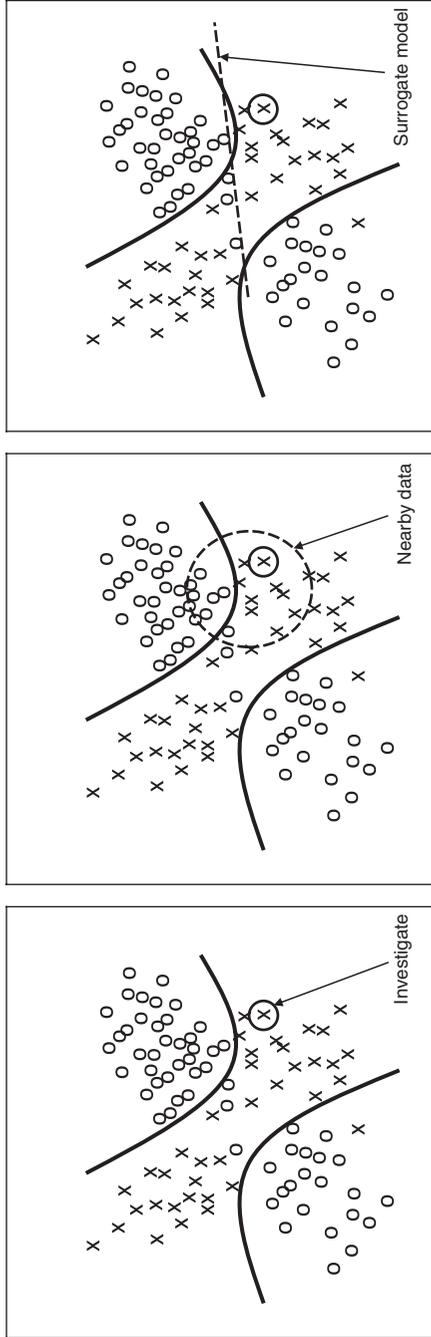


Figure 10.2 An example of how surrogate models can help with interpretability.

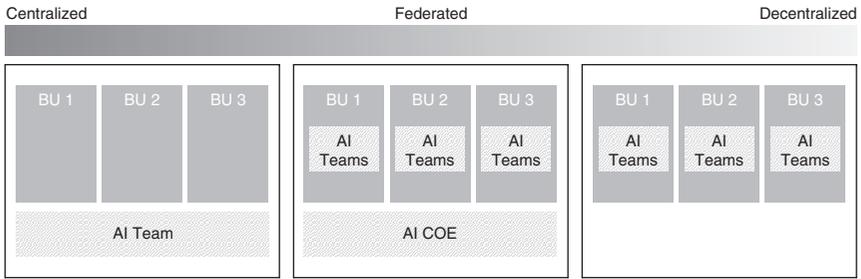


Figure 11.1 Centralized, decentralized, and federated operating models for AI.

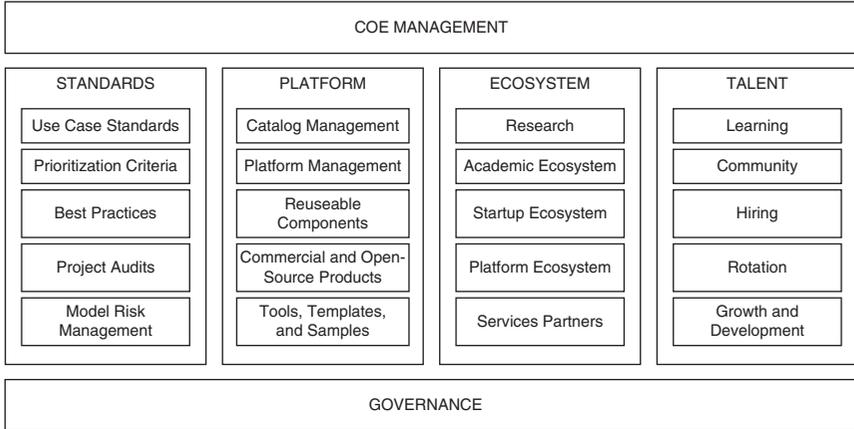


Figure 11.2 Key functions within an AI center of excellence.

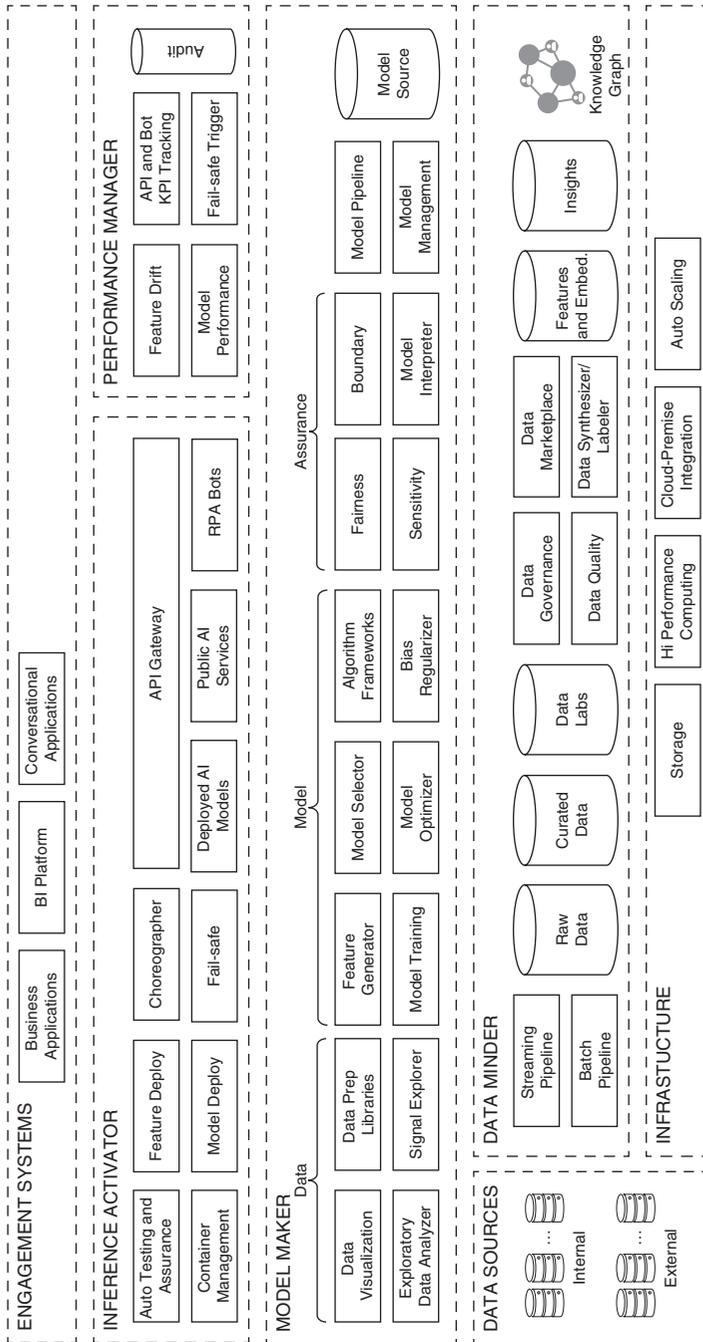


Figure 12.1 Architecture components for an AI platform.

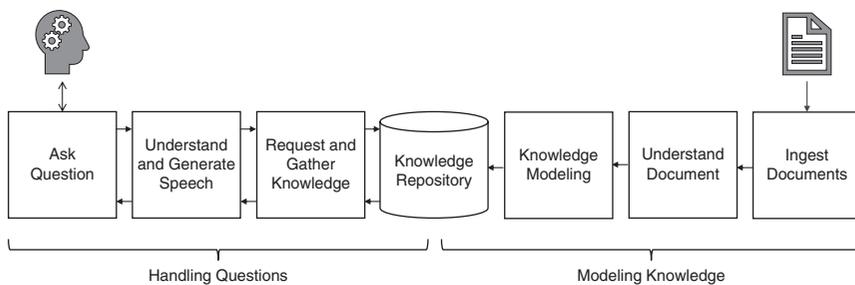


Figure 12.2 Question-and-answer systems built on knowledge modeling.

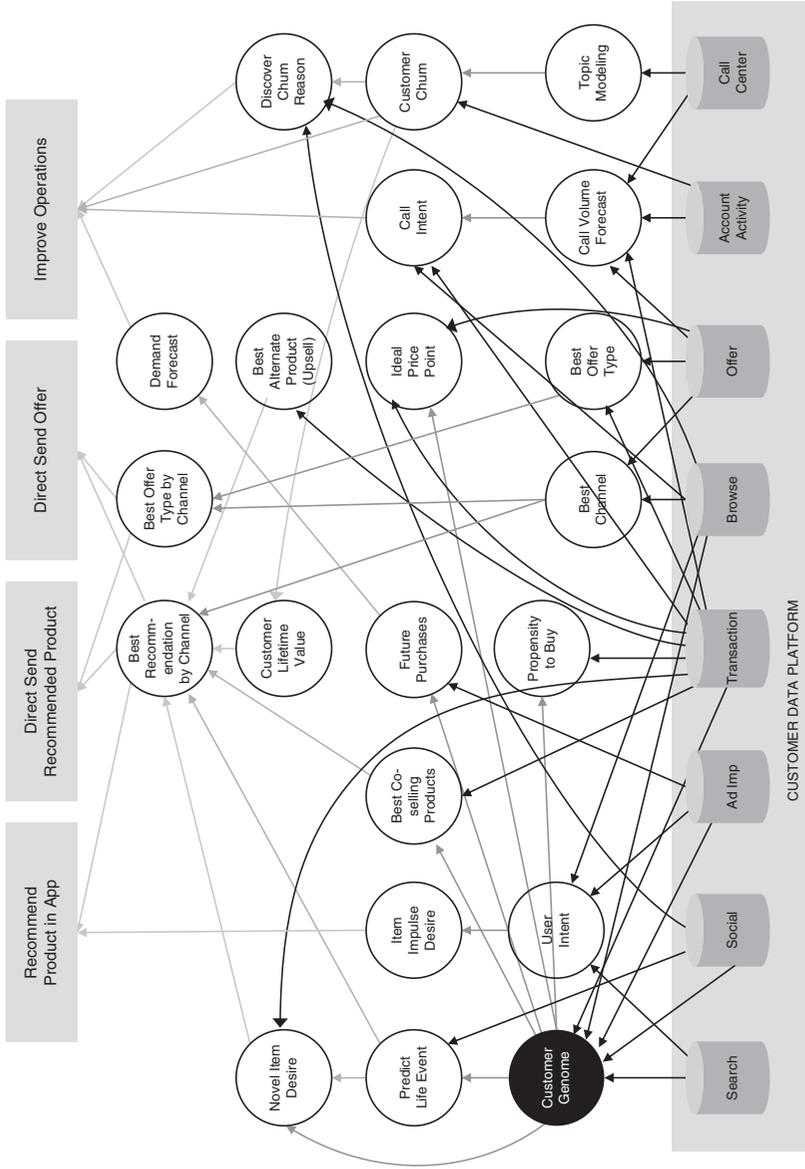


Figure 12.3 Leveraging multiple models for hyperpersonalization.

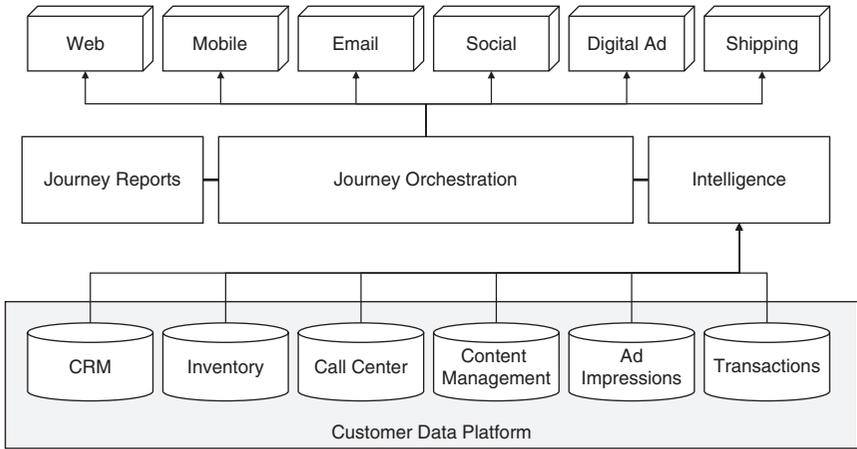


Figure 12.4 Orchestrating personalization interactions.

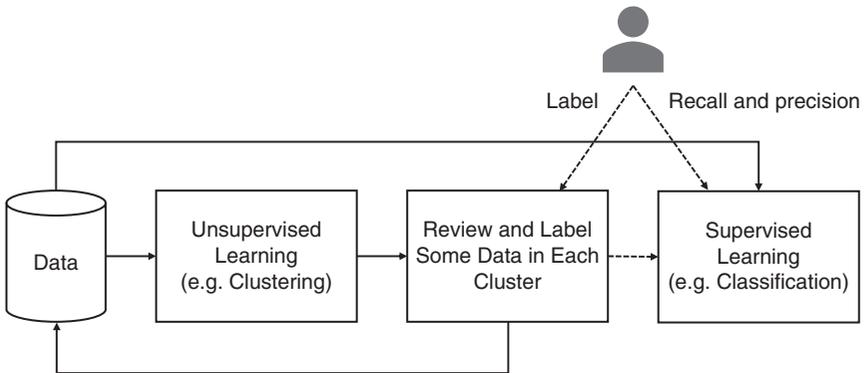


Figure 12.5 Activities for anomaly detection.

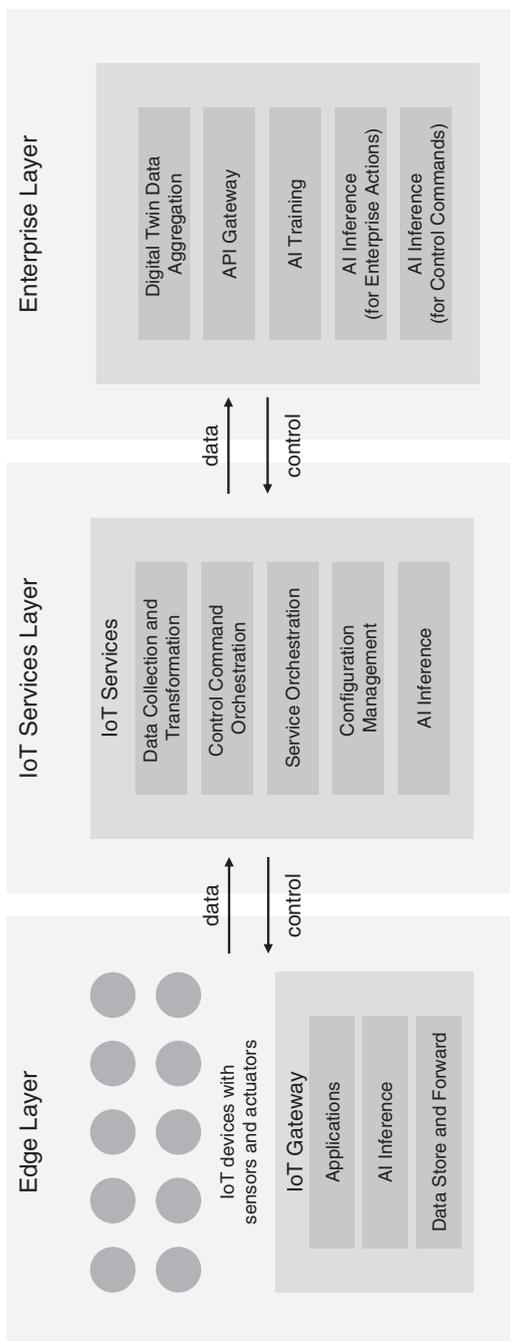


Figure 12.6 Interaction pattern for IoT and edge devices.

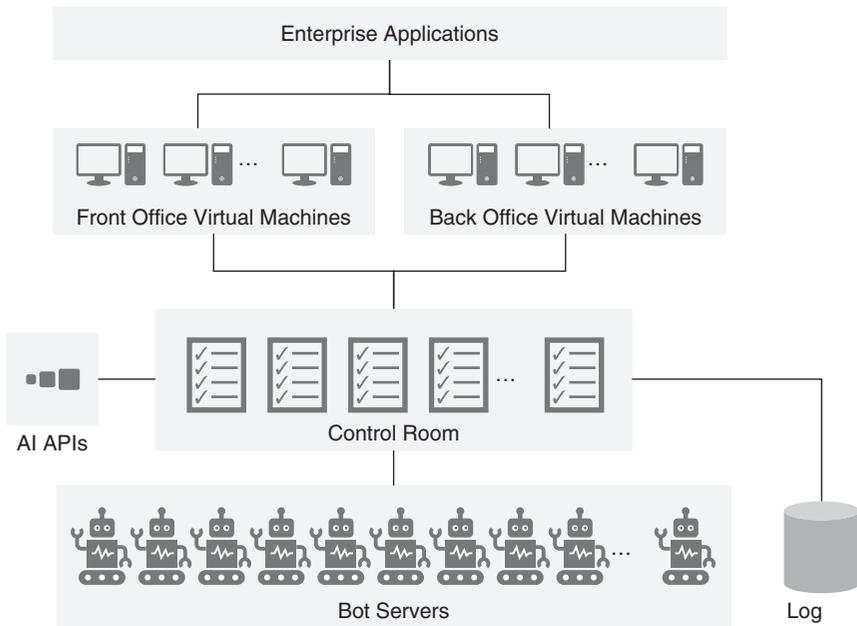


Figure 12.7 RPA-based digital workforce architecture.

```
1 # import math and data libraries
2 import pandas as pd
3 import numpy as np
4 from scipy.stats import uniform, randint
5
6 # import visualization libraries
7 import matplotlib.pyplot as plt
8 import seaborn as sns
9 from mpl_toolkits.mplot3d import Axes3D
10
11 import missingno as msno
12
13 from imblearn.over_sampling import SMOTE
14
15 # import sklearn machine learning libraries
16 from sklearn.preprocessing import LabelBinarizer, label_binarize, Imputer, \
17     LabelEncoder, OneHotEncoder, StandardScaler
18
19 from sklearn.compose import ColumnTransformer
20 from sklearn.pipeline import Pipeline
21 from sklearn.impute import SimpleImputer
22
23 from sklearn.model_selection import train_test_split, cross_val_score, \
24     GridSearchCV, KFold, StratifiedKFold, RandomizedSearchCV
25
26 # import the necessary model types
27 from sklearn.linear_model import LogisticRegression, LinearRegression
28 from sklearn.naive_bayes import GaussianNB
29 from sklearn.svm import LinearSVC
30 from sklearn.ensemble import RandomForestClassifier
31 import xgboost as xgb
32
33 # import model performance tools
34 from sklearn import metrics
35 from sklearn.metrics import precision_recall_curve, roc_curve, auc, \
36     accuracy_score, make_scorer, recall_score, \
37     precision_score, confusion_matrix
```

Figure 13.1 Importing relevant libraries that will be used.

```
1 # set the folder and file names from where you want to get data
2 folderName = 'gdrive/My Drive/Colab Notebooks/Data/'
3 fileName = 'customer_churn.csv'
4
5 # create dataframe and read file into dataframe
6 imp_data = pd.read_csv(folderName + fileName)
7 imp_data.shape
```

(3333, 22)

Figure 13.2 Importing the data for customer churn.

```
1 imp_data.head()
```

	CUSTOMER_ID	STATE	AREA_CODE	PHONE_NUMBER	ACCOUNT_LENGTH	INTL_PLAN	VMAII
0	10001	KS	415	382-4657	128	no	
1	10002	OH	415	371-7191	107	no	
2	10003	NJ	415	358-1921	137	no	
3	10004	OH	408	375-9999	84	yes	
4	10005	OK	415	330-6626	75	yes	

5 rows x 22 columns

Figure 13.3 Looking at the top few rows of the data.

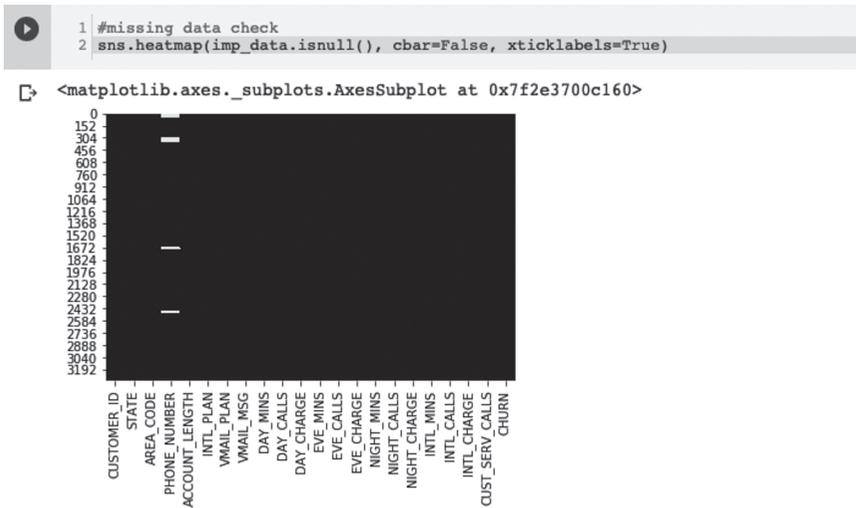


Figure 13.4 Heatmap of missing value. If there were any, they would show as a white bar for that row and column.

```

1 # drop features that are low impact
2 proc_data = imp_data.drop(columns=['AREA_CODE', 'PHONE_NUMBER'])
3 proc_data.shape
4
5 # transforming categorical data to numerical values
6 target_features = ['INTL_PLAN', 'VMAIL_PLAN', 'CHURN']
7 for i, target_feature in enumerate(target_features):
8     print(target_feature + " : ", proc_data[target_feature].unique())
9
10 # use encoder and transform
11 encoder = LabelEncoder()
12 for i, target_feature in enumerate(target_features):
13     encoded_values = encoder.fit_transform(proc_data[target_feature].values)
14     proc_data[target_feature] = pd.Series(encoded_values, index=imp_data.index)
15     # proc_data[target_feature] = proc_data[target_feature].astype('float64')
16     print(target_feature + " : ", proc_data[target_feature].unique())

```

↳ INTL_PLAN : ['no' 'yes']
VMAIL_PLAN : ['yes' 'no']
CHURN : ['False.' 'True.']
INTL_PLAN : [0 1]
VMAIL_PLAN : [1 0]
CHURN : [0 1]

Figure 13.5 Transforming categorical text data to numerical values.

```

1 # one hot encode categorical values that have more than 2 categories
2
3 proc_data = pd.get_dummies(proc_data, columns=['STATE'])
4 proc_data.shape

```

↳ (3333, 71)

Figure 13.6 One-hot encoding of US states.

```

1 # look at distribution of numeric data sets
2 col_names = ['ACCOUNT_LENGTH', 'VMAIL_MSG', 'DAY_MINS', 'DAY_CALLS', \
3             'DAY_CHARGE', 'EVE_MINS', 'EVE_CALLS', 'EVE_CHARGE', \
4             'NIGHT_MINS', 'NIGHT_CALLS', 'NIGHT_CHARGE', 'INTL_MINS', \
5             'INTL_CALLS', 'INTL_CHARGE', 'CUST_SERV_CALLS']
6
7 fig, axs = plt.subplots(5,3, figsize=(14,17))
8 for i, col_val in enumerate(col_names):
9     sns.distplot(proc_data[col_val], hist=True, ax=axs.flat[i])
10    axs.flat[i].set_xlabel(col_val, fontsize=8)
11    #axs.flat[i].set_ylabel('Count', fontsize=8)

```

Figure 13.7 Plotting frequency of datasets.

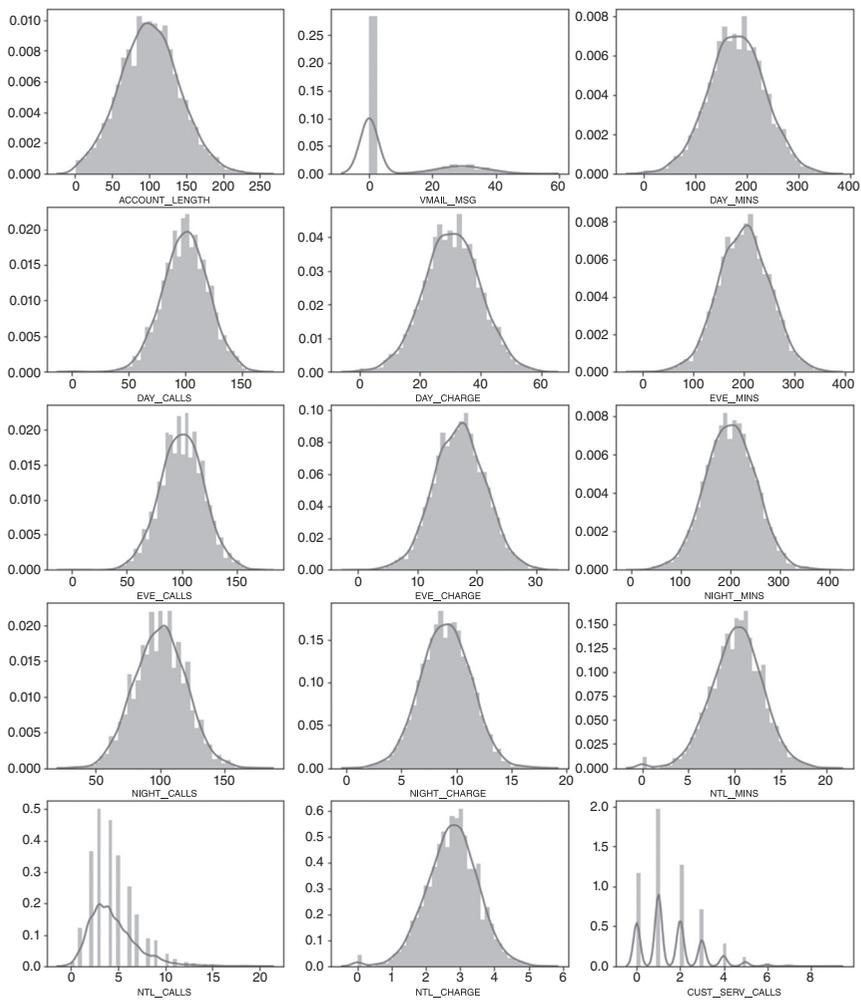


Figure 13.8 Frequency distribution of data of some of the columns.

```

1 # Create data subset for visualization
2 states = ['CUSTOMER_ID', 'STATE_AK', 'STATE_AL', 'STATE_AR', 'STATE_AZ', \
3           'STATE_CA', 'STATE_CO', 'STATE_CT', 'STATE_DC', 'STATE_DE', 'STATE_FL', \
4           'STATE_GA', 'STATE_HI', 'STATE_IA', 'STATE_ID', 'STATE_IL', 'STATE_IN', \
5           'STATE_KS', 'STATE_KY', 'STATE_LA', 'STATE_MA', 'STATE_MD', 'STATE_ME', \
6           'STATE_MI', 'STATE_MN', 'STATE_MO', 'STATE_MS', 'STATE_MT', 'STATE_NC', \
7           'STATE_ND', 'STATE_NE', 'STATE_NH', 'STATE_NJ', 'STATE_NM', 'STATE_NV', \
8           'STATE_NY', 'STATE_OH', 'STATE_OK', 'STATE_OR', 'STATE_PA', 'STATE_RI', \
9           'STATE_SC', 'STATE_SD', 'STATE_TN', 'STATE_TX', 'STATE_UT', 'STATE_VA', \
10          'STATE_VT', 'STATE_WA', 'STATE_WI', 'STATE_WV', 'STATE_WY']
11 viz_data = proc_data.drop(columns=states)
12 viz_data.shape
13
14 # Compute the correlation matrix
15 corr = viz_data.corr()
16
17 # Generate a mask for the upper triangle
18 mask = np.zeros_like(corr, dtype=np.bool)
19 mask[np.triu_indices_from(mask)] = True
20
21
22 # Set up the matplotlib figure
23 f, ax = plt.subplots(figsize=(11, 9))
24
25 # Generate a custom diverging colormap
26 cmap = sns.diverging_palette(220, 10, as_cmap=True)
27
28 # Draw the heatmap with the mask and correct aspect ratio
29 sns.heatmap(corr, mask=mask, cmap=cmap, vmax=.3, center=0,
30            square=True, linewidths=.5, cbar_kws={"shrink": .5})

```

↳ <matplotlib.axes._subplots.AxesSubplot at 0x7f2e37da4198>

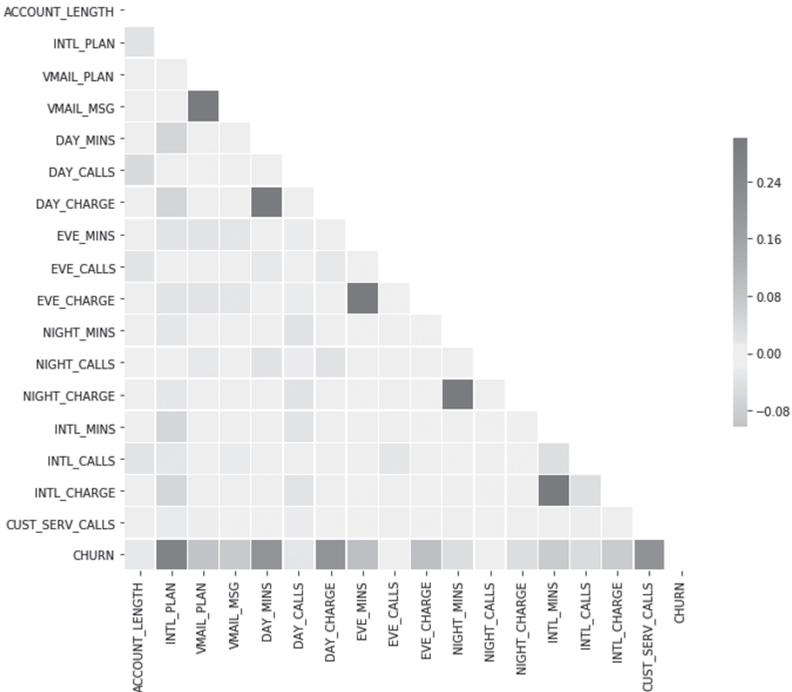


Figure 13.9 Heatmap of the correlations of some of the key columns with each other.

```

1 # exploring for outliers
2 col_names = ['ACCOUNT_LENGTH', 'DAY_MINS', 'DAY_CALLS', 'DAY_CHARGE', \
3             'EVE_MINS', 'EVE_CALLS', 'EVE_CHARGE', 'NIGHT_MINS', \
4             'NIGHT_CALLS', 'NIGHT_CHARGE']
5
6 fig, ax = plt.subplots(1, 10, figsize=(11,5))
7
8 for i, col_val in enumerate(col_names):
9     sns.boxplot(y=proc_data[col_val], ax=ax[i])
10    ax[i].set_ylabel('')
11    ax[i].set_xlabel(col_val, fontsize=8)

```

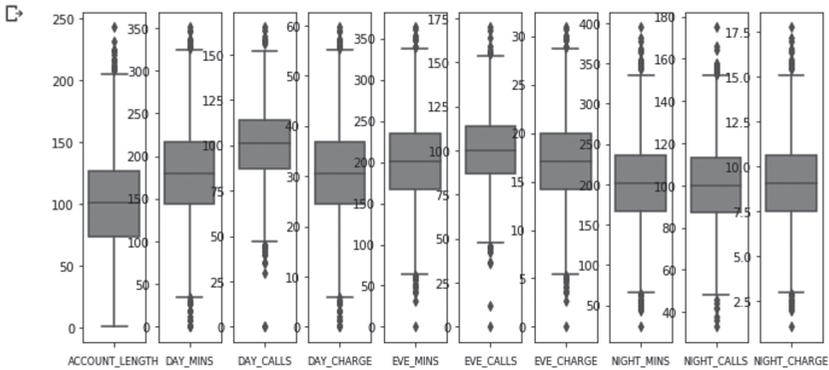


Figure 13.10 Looking for outliers.

```
1 # look at distribution of categorical data sets
2 num_col_names = ['INTL_PLAN', 'VMAIL_PLAN', 'CHURN']
3
4
5 fig, ax = plt.subplots(1, len(num_col_names), figsize=(11,6))
6 for i, col_val in enumerate(num_col_names):
7     sns.countplot(proc_data[col_val], ax=ax[i])
```

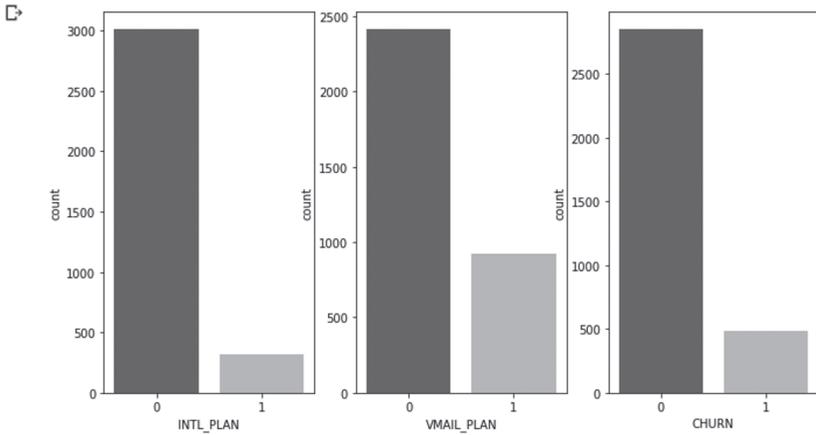


Figure 13.11 Imbalance in label or target data.

```

1 # scaling the features
2 scaler = StandardScaler()
3 scale_cols = ['ACCOUNT_LENGTH', 'DAY_MINS', 'DAY_CALLS', 'DAY_CHARGE', \
4             'EVE_MINS', 'EVE_CALLS', 'EVE_CHARGE', 'NIGHT_MINS', \
5             'NIGHT_CALLS', 'NIGHT_CHARGE',
6             'VMAIL_MSG', 'INTL_MINS', 'INTL_CALLS', 'INTL_CHARGE']
7 scaled_data = scaler.fit_transform(proc_data[scale_cols])
8 scaled_data = pd.DataFrame(scaled_data, columns=scale_cols)
9 scaled_full_data = proc_data.drop(scale_cols, axis=1)
10 scaled_full_data = pd.concat([scaled_full_data, scaled_data], \
11                             axis=1, sort=False)
12 scaled_full_data.shape
13
14 fig, (ax1, ax2) = plt.subplots(ncols=2, figsize=(12, 5))
15
16 ax1.set_title('Before Scaling')
17 sns.kdeplot(proc_data['ACCOUNT_LENGTH'], ax=ax1)
18 sns.kdeplot(proc_data['DAY_MINS'], ax=ax1)
19 sns.kdeplot(proc_data['DAY_CALLS'], ax=ax1)
20 sns.kdeplot(proc_data['DAY_CHARGE'], ax=ax1)
21 sns.kdeplot(proc_data['EVE_MINS'], ax=ax1)
22 sns.kdeplot(proc_data['EVE_CALLS'], ax=ax1)
23 sns.kdeplot(proc_data['EVE_CHARGE'], ax=ax1)
24 sns.kdeplot(proc_data['NIGHT_MINS'], ax=ax1)
25 sns.kdeplot(proc_data['NIGHT_CALLS'], ax=ax1)
26 sns.kdeplot(proc_data['NIGHT_CHARGE'], ax=ax1)
27
28 ax2.set_title('After Standard Scaler')
29 sns.kdeplot(scaled_data['ACCOUNT_LENGTH'], ax=ax2)
30 sns.kdeplot(scaled_data['DAY_MINS'], ax=ax2)
31 sns.kdeplot(scaled_data['DAY_CALLS'], ax=ax2)
32 sns.kdeplot(scaled_data['DAY_CHARGE'], ax=ax2)
33 sns.kdeplot(scaled_data['EVE_MINS'], ax=ax2)
34 sns.kdeplot(scaled_data['EVE_CALLS'], ax=ax2)
35 sns.kdeplot(scaled_data['EVE_CHARGE'], ax=ax2)
36 sns.kdeplot(scaled_data['NIGHT_MINS'], ax=ax2)
37 sns.kdeplot(scaled_data['NIGHT_CALLS'], ax=ax2)
38 sns.kdeplot(scaled_data['NIGHT_CHARGE'], ax=ax2)
39
40 plt.show()

```

Figure 13.12 Scaling the relevant data columns.

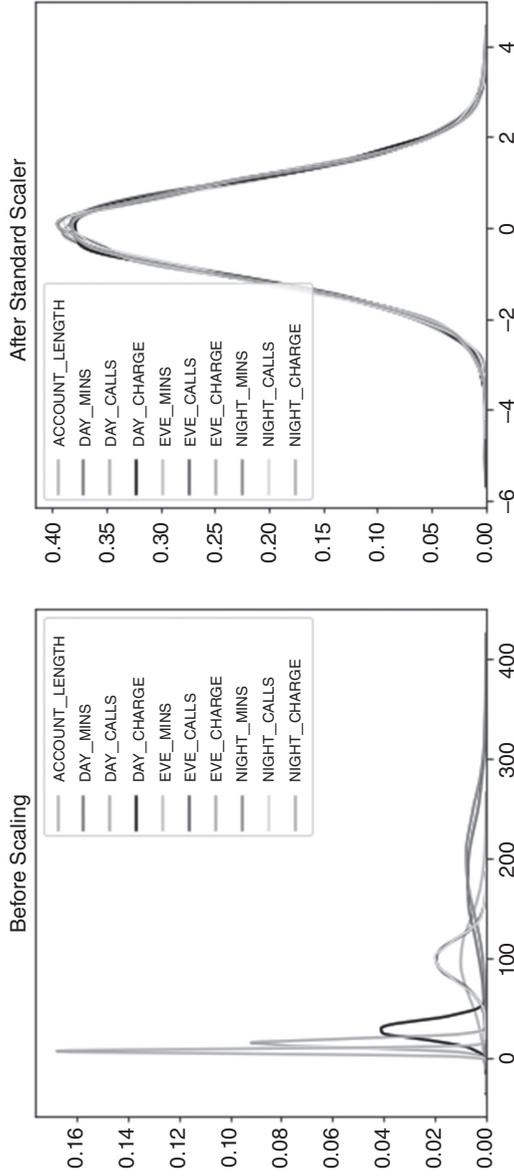


Figure 13.13 Visualizing the data distribution before scaling (left) and after scaling (right).

```

1 # remove columns with higher correlations, etc.
2 scaled_full_data['TOTAL_CHARGE'] = scaled_full_data['DAY_CHARGE'] + \
3   scaled_full_data['EVE_CHARGE'] + scaled_full_data['NIGHT_CHARGE'] + \
4   scaled_full_data['INTL_CHARGE']
5 scaled_full_data = scaled_full_data.drop(['DAY_CHARGE', 'EVE_CHARGE', \
6   'NIGHT_CHARGE', 'INTL_CHARGE'], axis = 1)

```

Figure 13.14 Dropping individual charge columns and adding the total charge column.

```

1 # churn by state, not using the one hot encoding
2 churn_by_state = pd.crosstab(imp_data.STATE, imp_data.CHURN, normalize='index')
3 churn_by_state = churn_by_state.sort_values(by='True.')
4 churn_by_state["True."].plot.bar(title="Churn by State", figsize=(11,3))

```

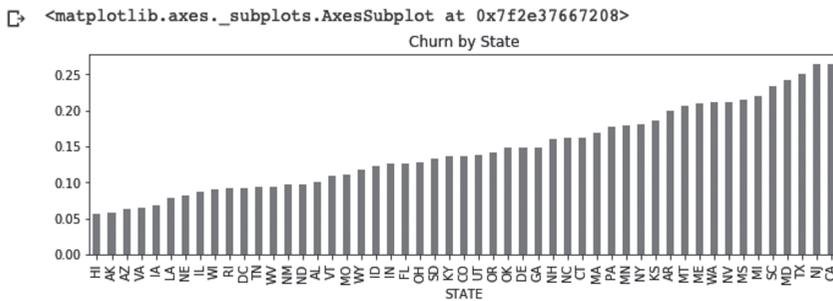


Figure 13.15 Analyzing churn rate by state.

```

1 # split the features from the target variable
2
3 sourcevars = scaled_full_data.drop(['CHURN', 'CUSTOMER_ID'], axis=1)
4 targetvar = scaled_full_data['CHURN']
5
6 # split the training and validation datasets
7
8 xTrain, xTest, yTrain, yTest = train_test_split(sourcevars, targetvar, \
9   test_size = 0.25, random_state = 0)
10
11 sourcevars.shape, targetvar.shape

```

((3333, 65), (3333,))

Figure 13.16 Splitting data for training and testing in the ratio of 75:25.

```

1 # try classification models
2 model = LogisticRegression(solver = 'lbfgs')
3
4 # train the algorithm on training data and predict using the testing data
5 model.fit(xTrain, yTrain)
6 predictions = model.predict(xTest)
7 print("Accuracy : ",accuracy_score(yTest, predictions, normalize = True))

```

Accuracy : 0.8477218225419664

Figure 13.17 Set up a logistic regression model for binary classification.

```

1 # this is the accuracy if you assume NO customers will churn
2 1 - yTest.mean()

```

0.8621103117505995

Figure 13.18 Percentage of customers that did not churn in the validation dataset.

```

1 # print(metrics.confusion_matrix(yTest, predictions))
2 print(pd.DataFrame(confusion_matrix(yTest, predictions),
3                    columns=['pred_no_churn', 'pred_churn'],
4                    index=['actual_no_churn', 'actual_churn']))

```

	pred_no_churn	pred_churn
actual_no_churn	687	32
actual_churn	95	20

```

[173] 1 # look at performance metrics
      2 print(metrics.classification_report(yTest, predictions))

```

	precision	recall	f1-score	support
0	0.88	0.96	0.92	719
1	0.38	0.17	0.24	115
accuracy			0.85	834
macro avg	0.63	0.56	0.58	834
weighted avg	0.81	0.85	0.82	834

Figure 13.19 Looking at the confusion matrix and precision, recall, and F1 score.

```
1 # create ROC curve
2 plt.style.use('ggplot')
3 y_predict_probabilities = model.predict_proba(xTest)[:,-1]
4
5 fpr, tpr, _ = roc_curve(yTest, y_predict_probabilities)
6 roc_auc = auc(fpr, tpr)
7
8 plt.figure()
9 plt.plot(fpr, tpr, color='darkorange', lw=2, \
10         label='ROC curve (area = %0.2f)' % roc_auc)
11 plt.plot([0, 1], [0, 1], color='navy', lw=2, linestyle='--')
12 plt.xlim([0.0, 1.0])
13 plt.ylim([0.0, 1.05])
14 plt.xlabel('False Positive Rate')
15 plt.ylabel('True Positive Rate')
16 plt.title('ROC Curve')
17 plt.legend(loc="lower right")
18 plt.show()
```

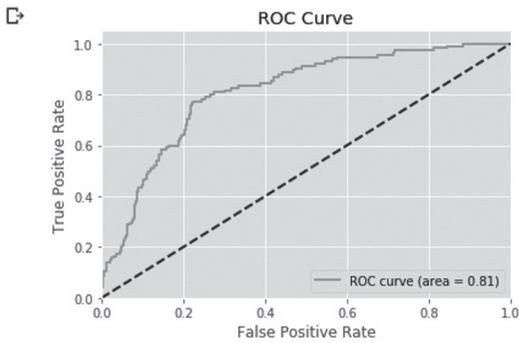


Figure 13.20 Receiver operating characteristic (ROC) curve and area under the curve (AUC).

```

1 # data augmentation
2
3 sm = SMOTE(random_state = 2)
4 xTrainBal, yTrainBal = sm.fit_sample(xTrain, yTrain.ravel())
5 predictions = model.fit(xTrainBal, yTrainBal).predict(xTest)

```

```

[183] 1 print("Accuracy : ",accuracy_score(yTest, predictions, normalize = True))

```

↳ Accuracy : 0.7386091127098321

```

[184] 1 # print(metrics.confusion_matrix(yTest, predictions))
2 print(pd.DataFrame(confusion_matrix(yTest, predictions),
3                    columns=['pred_no_churn', 'pred_churn'],
4                    index=['actual_no_churn', 'actual_churn']))

```

↳

	pred_no_churn	pred_churn
actual_no_churn	530	189
actual_churn	29	86

```

[185] 1 # look at performance metrics
2 print(metrics.classification_report(yTest, predictions))

```

↳

	precision	recall	f1-score	support
0	0.95	0.74	0.83	719
1	0.31	0.75	0.44	115
accuracy			0.74	834
macro avg	0.63	0.74	0.64	834
weighted avg	0.86	0.74	0.78	834

Figure 13.21 Augmenting the minority data.

```

1 # try classification models
2 # model = LogisticRegression(solver = 'lbfgs')
3 model = xgb.XGBClassifier(objective="binary:logistic", random_state=42)
4
5 # train the algorithm on training data and predict using the testing data
6 model.fit(xTrain, yTrain)
7 predictions = model.predict(xTest)
8 print("Accuracy : ",accuracy_score(yTest, predictions, normalize = True))

```

↳ Accuracy : 0.960431654676259

```

[189] 1 # print(metrics.confusion_matrix(yTest, predictions))
2 print(pd.DataFrame(confusion_matrix(yTest, predictions),
3                    columns=['pred_no_churn', 'pred_churn'],
4                    index=['actual_no_churn', 'actual_churn']))

```

	pred_no_churn	pred_churn
actual_no_churn	710	9
actual_churn	24	91

```

[190] 1 # look at performance metrics
2 print(metrics.classification_report(yTest, predictions))

```

	precision	recall	f1-score	support
0	0.97	0.99	0.98	719
1	0.91	0.79	0.85	115
accuracy			0.96	834
macro avg	0.94	0.89	0.91	834
weighted avg	0.96	0.96	0.96	834

Figure 13.22 Trying a different algorithm – only lines 2 and 3 in the first block have been changed to select a different model.

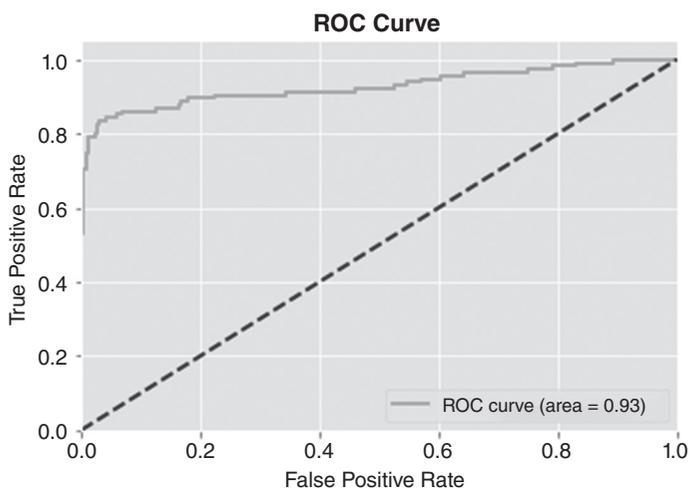


Figure 13.23 ROC curve and AUC using XGBoost.

```
1 # explore feature importance
2 fig, ax = plt.subplots(1,1,figsize=(9,6))
3 xgb.plot_importance(model, max_num_features=10, ax=ax)
4 plt.show()
```

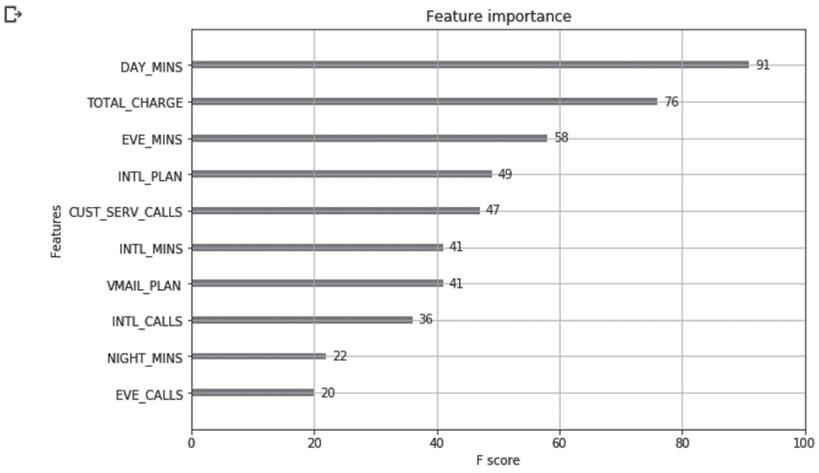


Figure 13.24 Feature importance for the top 10 features in the model.