

▼ Data Pre-Processing

David Gerbing
The School of Business
Portland State University
gerbing@pdx.edu

- [1. Preliminaries](#)
 - [1.1 Packages](#)
 - [1.2 Read](#)
- [2. Create Dummy Variables](#)
- [3. Missing Data](#)
 - [3.1 Assess Amount of Missing Data](#)
 - [3.2 Show Rows with Missing Data](#)
 - [3.3 Delete rows with Missing Data](#)
 - [3.4 Impute Missing Data](#)
- [4. Search for Outliers](#)
- [5. Transform Variables to Similar Scale](#)
 - [5.1 Min-Max Scaling](#)
 - [5.1.1 Apply to Original Data](#)
 - [5.1.2 Apply to New Data](#)
 - [5.2 Standardization Scaling](#)
 - [5.3 Robust Scaling](#)

▼ Preliminaries

▼ Packages

```
from datetime import datetime as dt
now = dt.now()
print("Analysis on", now.strftime("%Y-%m-%d"), ", at", now.strftime("%H:%M"))
Analysis on 2021-06-25 at 01:11
```

```
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
import seaborn as sns
```

▼ Read

```
#d = pd.read_excel('data/employee.xlsx')
d = pd.read_excel('http://lessstats.com/data/employee.xlsx')
```

The data values in the Name column are not values per se to analyze, but instead serve as row identifiers, ID's. As such, replace the default integer row labels with the values of the column Name. Do so with the `set_index()` function.

```
d = d.set_index('Name')
d.head()
```

	Years	Gender	Dept	Salary	JobSat	Plan	Pre	Post
Name								
Ritchie, Darnell	7.0	M	ADMN	53788.26	med	1	82	92
Wu, James	NaN	M	SALE	94494.58	low	1	62	74
Hoang, Binh	15.0	M	SALE	111074.86	low	3	96	97
Jones, Alissa	5.0	F	NaN	53772.58	NaN	1	65	62
Downs, Deborah	7.0	F	FINC	57139.90	high	2	90	86

▼ Create Dummy Variables

designateMachine learning procedures cannot directly process categorical variables with non-numeric values. For example, consider a data set with two values of `Gender`, coded as M and F. The variable `Gender` with this non-numerical coding cannot be entered into a machine learning analysis, which requires numerical variables only.

Categorical variables with non-numerical values, however, can be converted to numerical representations. Many such conversions are possible. Here we consider the most widely used conversion.

Dummy Variable: A numerically encoded variable for each level of a categorical variable, with a value of 1 if the level is present and 0 if not.

The `Gender` variable becomes two dummy variables, `Gender_F` and `Gender_M`. For example, if a person's `Gender` is listed as F, then `Gender_F` is 0 and `Gender_M` is 1.

Pandas provides the function `get_dummies()` to convert a categorical variable to a corresponding set of dummy values, one for each category. The parameter `columns` designates the variables to be converted.

One adjustment is needed. If you know the value of `Gender_F` for an individual is 1, then you also know that `Gender_M` is 0. So the value of either one of two dummy variables implies the value of the other. To avoid redundancy, in general, for k levels of the categorical variables, the number of dummy variables retained in the analysis is $k - 1$. For two levels of `Gender`, arbitrarily retain $2 - 1 = 1$ of the dummy variables in the analysis.

With `get_dummies()`, drop the first dummy variable with the `drop_first` parameter set to `True`. Alphabetically, F comes before M, so in the following analysis, `Gender_F` is dropped. The original `Gender` variable is replaced with `Gender_M`.

For `JobSat` with three levels – High, Low, and Med – create three dummy variables, each corresponding to the one of the three values. `JobSat` is then replaced by two dummy variables for the Low and Med values. For example, if you know the values of `JobSat_low` and `JobSat_med` are both 0, then you know that the value of `JobSat_high` is 1. Knowing two values implies the third, so retain only two dummy variables for `JobSat` in the analysis.

```
d = pd.get_dummies(d, columns=['Gender', 'JobSat'], drop_first=True)
d.head()
```

	Years	Dept	Salary	Plan	Pre	Post	Gender_M	JobSat_low	JobSat_med
Name									
Ritchie, Darnell	7.0	ADMN	53788.26	1	82	92	1	0	1
Wu, James	NaN	SALE	94494.58	1	62	74	1	1	0
Hoang, Binh	15.0	SALE	111074.86	3	96	97	1	1	0
Jones, Alissa	5.0	NaN	53772.58	1	65	62	0	0	0
Downs, Deborah	7.0	FINC	57139.90	2	90	86	0	0	0

Usually which particular dummy variable is dropped for each categorical variable is irrelevant. If, however, it is desired to drop a dummy variable other than the first, then run `get_dummies()` without the `drop_first` parameter and manually drop the specified dummy variable from the data frame.

▼ Missing Data

Machine learning functions generally do not work in the presence of missing data. Before machine learning analysis, examine the data for missing data and adjust accordingly, either delete the row or column or impute the value.

A missing data value is indicated by the notation `NaN`, an abbreviation for Not a Number. Sometimes functions or discussion of missing data refer to missing data as `na`, which means Not Available.

Here James Wu has a missing value for the number of years he worked at the company. Data values for James Wu occupy the second row of data, identified by row index 1. The row definition of 1:2 (confusingly) also refers to the second row. However, specifying the row as a range results in the output's more visually appealing horizontal placement.

```
d.iloc[1:2, 0:5]
```

	Years	Dept	Salary	Plan	Pre
Name					
Wu, James	NaN	SALE	94494.58	1	62

The `isna()` function indicates if a data value is missing. Follow with the `sum()` function to sum the number of missing values for a variable, here all variables in the data frame because no specific variable is specified. Follow with a second `sum()` function to sum the sums, that is, the total number of missing values in the entire data frame.

```
print(d.isna().sum())
print('\nTotal Missing:', d.isna().sum().sum())
```

Years 1
Dept 1
Salary 0
Plan 0
Pre 0
Post 0
Gender_M 0
JobSat_low 0
JobSat_med 0
dtype: int64

Total Missing: 2

As a programming note, without using the `print()` function, the last row of code in a Jupyter cell that specifies output generates the default output. If there is more than a single line of code that generates output, or if customization of the output is desired, such as adding a descriptive label, then invoke `print()`, as in this example.

▼ Show Rows with Missing Data

The code for viewing all rows of missing data begins with the `isna()` function, which returns `True` if a data value is missing. The `any()` function evaluates the data frame column-by-column and then returns `True` if there are any `True` values in the corresponding row. Putting the expression within `d[::]` selects only the rows with `True`, that is, with missing according to `isna()`.

```
d[d.isna().any(axis='columns')]
```

	Years	Dept	Salary	Plan	Pre	Post	Gender_M	JobSat_low	JobSat_med
Name									
Wu, James	NaN	SALE	94494.58	1	62	74	1	1	0

The `isna()` function indicates if a data value is missing. Follow with the `sum()` function to sum the number of missing values for a variable, here all variables in the data frame because no specific variable is specified. Follow with a second `sum()` function to sum the sums, that is, the total number of missing values in the entire data frame.

```
print(d.isna().sum())
print('\nTotal Missing:', d.isna().sum().sum())
```

Years 1
Dept 1
Salary 0
Plan 0
Pre 0
Post 0
Gender_M 0
JobSat_low 0
JobSat_med 0
dtype: int64

Total Missing: 2

As a programming note, without using the `print()` function, the last row of code in a Jupyter cell that specifies output generates the default output. If there is more than a single line of code that generates output, or if customization of the output is desired, such as adding a descriptive label, then invoke `print()`, as in this example.

▼ Delete rows with Missing Data

The simplest method to address missing data deletes a row if it contains any missing data, what is called case deletion, or list-wise deletion. The `dropna()` function deletes rows with missing data from d. It is also possible to apply the function to columns with parameter `axis`, which indicates if the analysis applies to rows or columns. Often in Python coding people use 0 instead of the more descriptive 'rows'.

The process in this example removes the three rows with missing data, from 37 rows to 34 rows.

```
d.shape
```

```
(37, 9)
```

```
d = d.dropna()
d.shape
```

```
(35, 9)
```

The problem with dropping rows that contain missing data is that for some data sets much or most of the data can be deleted. Appropriate if many data values in an entire row are missing, but perhaps not if just one missing data value across data for many variables. Or, sometimes a single variable may contain many missing values, so better to delete the bad variable than delete so many corresponding rows of data (cases).

To illustrate, re-display the variable names and the first five rows of data. In the original data frame, James Wu is missing Years worked for the company, and Alissa Jones is missing the Dept worked in as well and the Job Satisfaction rating. Both rows of data are now deleted from the revised data frame.

```
d.head()
```

	Years	Dept	Salary	Plan	Pre	Post	Gender_M	JobSat_low	JobSat_med
Name									
Ritchie, Darnell	7.0	ADMN	53788.26	1	82	92	1	0	1
Hoang, Binh	15.0	SALE	111074.86	3	96	97	1	1	0
Downs, Deborah	7.0	FINC	57139.90	2	90	86	0	0	0
Jones, Alissa	5.0	NaN	53772.58	1	65	62	0	0	0
Knox, Michael	18.0	MKTG	99062.66	3	81	84	1	0	1

Usually which particular dummy variable is dropped for each categorical variable is irrelevant. If, however, it is desired to drop a dummy variable other than the first, then run `get_dummies()` without the `drop_first` parameter and manually drop the specified dummy variable from the data frame.

▼ Create Dummy Variables

Machine learning functions generally do not work in the presence of missing data. Before machine learning analysis, examine the data for missing data and adjust accordingly, either delete the row or column or impute the value.

A missing data value is indicated by the notation `NaN`, an abbreviation for Not a Number. Sometimes functions or discussion of missing data refer to missing data as `na`, which means Not Available.

Here James Wu has a missing value for the number of years he worked at the company. Data values for James Wu occupy the second row of data, identified by row index 1. The row definition of 1:2 (confusingly) also refers to the second row. However, specifying the row as a range results in the output's more visually appealing horizontal placement.

```
d[d.isna().any(axis=1)]
```

	Years	Dept	Salary	Plan	Pre	Post	Gender_M	JobSat_low	JobSat_med
Name									
Wu, James	NaN	SALE	94494.58	1	62	74	1	1	0

The `isna()` function indicates if a data value is missing. Follow with the `sum()` function to sum the number of missing values for a variable, here all variables in the data frame because no specific variable is specified. Follow with a second `sum()` function to sum the sums, that is, the total number of missing values in the entire data frame.

```
print(d.isna().sum())
print('\nTotal Missing:', d.isna().sum().sum())
```

Years 1
Dept 1
Salary 0
Plan 0
Pre 0
Post 0
Gender_M 0
JobSat_low 0
JobSat_med 0
dtype: int64

Total Missing: 2

As a programming note, without using the `print()` function, the last row of code in a Jupyter cell that specifies output generates the default output. If there is more than a single line of code that generates output, or if customization of the output is desired, such as adding a descriptive label, then invoke `print()`, as in this example.

▼ Show Rows with Missing Data

The code for viewing all rows of missing data begins with the `isna()` function, which returns `True` if a data value is missing. The `any()` function evaluates the data frame column-by-column and then returns `True` if there are any `True` values in the corresponding row. Putting the expression within `d[::]` selects only the rows with `True`, that is, with missing according to `isna()`.

```
d[d.isna().any(axis='columns')]
```

	Years	Dept	Salary	Plan	Pre	Post	Gender_M	JobSat_low	JobSat_med
Name									
Wu, James	NaN	SALE	94494.58	1	62	74	1	1	0

The `isna()` function indicates if a data value is missing. Follow with the `sum()` function to sum the number of missing values for a variable, here all variables in the data frame because no specific variable is specified. Follow with a second `sum()` function to sum the sums, that is, the total number of missing values in the entire data frame.

```
print(d.isna().sum())
print('\nTotal Missing:', d.isna().sum().sum())
```

Years 1
Dept 1
Salary 0
Plan 0
Pre 0
Post 0
Gender_M 0
JobSat_low 0
JobSat_med 0
dtype: int64

Total Missing: 2

Standardization Scaling

Another linear transformation of the data values converts the original data values to *z*-scores, this one taught in all introductory stat courses.

Standard score: The number of standard deviations a data value is from the mean.

Write the transformation of the data values for variable *x* as:

$$z_i = \frac{x_i - m}{s}$$

To standardize, for each data value of a variable, for the *i*th row of data, subtract the mean of the data and divide by the standard deviation of the data. The result, a distribution of *z*-scores, has a mean of 0 and a standard deviation of 1.

For this linear transformation, set *a* = $-(m/s)$ and *b* = $1/s$, where *s* is the sample standard deviation and *m* is the sample mean.

The `StandardScaler` provides the computations for standardization of variables. If (not a requirement for standardization) a variable is normal, then most values will be within 2.5 or 3 standard deviations from the mean, that is, standard scores of less than 3.0 and greater than -3.0.

```
from sklearn.preprocessing import StandardScaler
```

```
s_scaler = preprocessing.StandardScaler()
```

Get the mean and standard deviation of each variable with the `fit()` function. Do the rescaling with the `transform()` function. Combine both with the `fit_transform()` function.

```
Xst = s_scaler.fit_transform(X)
Xst = pd.DataFrame(Xst, columns=['Years', 'Salary', 'Pre'])
Xst.head()
```

	Years	Salary	Pre
0	-0.443135	-0.926820	0.201793
1	0.966840	1.729495	1.407629
2	-0.443135	-0.771408	0.890842
3	-0.619382	-0.200977	1.752154
4	1.495581	1.172503	0.115662

The success of the transformation is shown by examining the mean and standard deviation of the three transformed, now standardized, variables.

```
round(Xst.mean(), 4)
```

	Years	Salary	Pre
0	0.0	0.0	0.0
1	1.0146	1.0146	1.0146
2	-0.0	-0.0	-0.0
3	-1.779224	-1.779224	-1.779224
4	1.0146	1.0146	1.0146

round(Xst.std(), 4)

	Years	Salary	Pre
0	1.0146	1.0146	1.0146
1	1.0146	1.0146	1.0146
2	1.0146	1.0146	1.0146
3	1.0146	1.0146	1.0146
4	1.0146	1.0146	1.0146

The range of the data roughly approximates that of normal data, though skewed right. In a perfectly normal distribution the standardized values would range from about -2.5 to 2.5.

```
Xst.min()
```

	Years	Salary	Pre
0	-1.500617	-1.282158	-1.779224
1	1.282158	1.811947	1.779224
2	-1.500617	-1.779224	-1.779224
3	-1.779224	-1.779224	-1.779224
4	1.282158	1.811947	1.779224

```
Xst.max()
```

	Years	Salary	Pre
0	2.553063	2.811947	1.752154
1	2.811947	2.811947	1.752154
2	2.553063	2.811947	1.752154
3	2.811947	2.811947	1.752154
4	2.811947	2.811947	1.752154

Can transform any data with the same values from `fit()`.

```
s_scaler.transform([[15.0, 111074.86, 96.0]])
array([[0.96684042, 1.72949472, 1.4076293]])
```

Can also transform manually. The computed values by which each data value is transformed is available from the computed `scale_` and `mean_` data structures. Each computed data structure contains one value for each variable in the data frame that is transformed.

```
print('mean:', s_scaler.mean_)
```

```
print('sd:', s_scaler.scale_)
```

```
mean: [19.51428571e+00 7.37762411e+04 7.96571429e+01]
sd: [5.67385698e+00 2.15661941e+04 1.16101996e+01]
```

To illustrate, compute the transformed value of *Salary* for the second row of data in *X*.

```
(111074.86 - s_scaler.mean_[1]) / s_scaler.scale_[1]
```

```
1.7294947196040262
```

```
Xst.head()
```

	Years	Salary	Pre
--	-------	--------	-----

	Name	Years	Salary	Pre
--	------	-------	--------	-----

0	Ritchie, Darnell	7.0	53788.26	82.0
1	Hoang, Binh	15.0	111074.86	96.0
2	Downs, Deborah	7.0	57139.90	90.0
3	Afshari, Anbar	6.0	69441.93	100.0
4	Knox, Michael	18.0	99062.66	81.0

Robust Scaling

Robust scaling resembles standardization, except it is more robust to the presence of outliers. The presence of outliers does not dramatically change the resulting scaled values as much as standardization in which an outlier can have a significant impact on the mean and an even bigger impact on increasing the size of the standard deviation (which depends on squared deviation scores).

Robust scaling accomplishes this robustness by replacing the mean in the standard score formula with the more robust median and the standard deviation with the more robust interquartile range. The median is the second quartile, and the IQR is the difference between the third and first quartiles. Unlike the mean and standard deviation, no matter how extreme a few values are in a distribution, the quartiles remain the same.

Robust scale score: The number of IQR's a data value is from the median.

Write the transformation of the data values for variable *x* as:

$$\text{robustscore} = \frac{x_i - \text{median}}{\text{IQR}}$$

That is, to do a robust scaling, for each data value of a variable, for the *i*th row of data, subtract the median of the data and divide by the IQR of the data.

```
from sklearn.preprocessing import RobustScaler
```

```
r_scaler = preprocessing.RobustScaler()
```

```
Xrb = r_scaler.fit_transform(X)
Xrb = pd.DataFrame(Xrb, columns=['Years', 'Salary', 'Pre'])
Xrb.head()
```

	Years	Salary	Pre
--	-------	--------	-----

	Name	Years	Salary	Pre
--	------	-------	--------	-----

0	-0.250	-0.554057	0.108108	
1	0.750	1.459988	0.864865	
2	-0.250	-0.436222	0.540541	
3	-0.375	-0.03715	1.081081	
4	1.125	1.037672	0.054054	

The specific characteristics of the transformed variables differ from standardization, but the general results remain. The means are somewhat close to 0. The standard deviations are less than 1 but certainly much closer to 1 than from the original distributions. The minimum and maximum values are less than the range of the standardized variables but roughly similar, again, especially compared to the original distributions.

```
round(Xrb.mean(), 4)
```

	Years	Salary	Pre
--	-------	--------	-----

	Name	Years	Salary	Pre
--	------	-------	--------	-----

0	Ritchie, Darnell	0.0643	0.1487	
1	Hoang, Binh	0.7693	0.6367	
2	Downs, Deborah	-0.0185	-0.1351	
3	Afshari, Anbar	-0.6367	-0.7196	
4	Knox, Michael	1.081081	0.6367	

```
round(Xrb.std(), 4)
```

	Years	Salary	Pre
--	-------	--------	-----

	Name	Years	Salary	Pre
--	------	-------	--------	-----

0	Ritchie, Darnell	0.7196	0.7693	
1	Hoang, Binh	0.6367	0.6367	
2	Downs, Deborah	-0.1351	-0.7196	
3	Afshari, Anbar	-0.7196	-0.6367	
4	Knox, Michael	1.081081	0.6367	

```
round(Xrb.min(), 4)
```

	Years	Salary	Pre
--	-------	--------	-----

	Name	Years	Salary	Pre
--	------	-------	--------	-----

0	Ritchie, Darnell	-0.250	-0.554057	
1	Hoang, Binh	0.750	1.459988	
2	Downs, Deborah	-0.250	-0.436222	
3	Afshari, Anbar	-0.375	-0.03715	
4	Knox, Michael	1.125	1.037672	

```
round(Xrb.max(), 4)
```

	Years	Salary	Pre
--	-------	--------	-----

	Name	Years	Salary	Pre
--	------	-------	--------	-----

0	Ritchie, Darnell	0.7196	0.7693	
1	Hoang, Binh	0.6367	0.6367	
2	Downs, Deborah	-0.1351	-0.7196	
3	Afshari, Anbar	-0.7196	-0.6367	
4	Knox, Michael	1.081081	0.6367	