


Lecture 19: Dealing with Missing Data

CS109A Introduction to Data Science

Pavlos Protopapas, Kevin Rader and Chris Tanner



Lecture Outline: Multiclass Classification

Dealing with Missingness

Types of Missingness

Imputation Methods

Dealing with Missingness

What is missing data? k-Nearest Neighbors

Often times when data is collected, there are some missing values apparent in the dataset. This leads to a few questions to consider:

- How does this show up in pandas?
- How does sklearn handle these NaNs?
- How does this effect our modeling?

What are the simplest ways to handle missing data?

Naively handling missingness

- Drop the observations that have any missing values.
 - Use `pd.DataFrame.dropna(axis=0)`
- Impute the mean/median (if quantitative) or most common class (if categorical) for all missing values.
 - Use `pd.DataFrame.fillna(value=x.mean())`

How do statsmodels and sklearn handle these NaNs?

What are some consequences in handling missingness in these fashions?

Missingness Indicator Variable

One simple way to handle missingness in a variable, X_j , is to impute a value (like 0 or \bar{X}_j), then create a new variable, $X_{j,miss}$, that indicates this observation had a missing value. If X_j is categorical then just impute 0.

Then include both $X_{j,miss}$ and X_j as predictors in any model.

Illustration is to the right.

| X_1 | X_2 | X_1^* | X_2^* | $X_{1,miss}$ | $X_{2,miss}$ |
|-------|-------|---------|---------|--------------|--------------|
| 10 | . | 10 | 0 | 0 | 1 |
| 5 | 1 | 5 | 1 | 0 | 0 |
| 21 | 0 | 21 | 0 | 0 | 0 |
| 15 | 0 | 15 | 0 | 0 | 0 |
| 16 | . | 16 | 0 | 0 | 1 |
| . | . | 0 | 0 | 1 | 1 |
| 21 | 1 | 21 | 1 | 0 | 0 |
| 12 | 0 | 12 | 0 | 0 | 0 |
| . | 1 | 0 | 1 | 1 | 0 |

Why use a Missingness Indicator Variable?

How does this missingness indicator variable improve the model?

Because the group of individuals with a missing entry may be systematically different than those with that variable measured. Treating them equivalently could lead to bias in quantifying relationships (the β 's) and underperform in prediction.

For example: imagine a survey that asks whether or not someone has ever recreationally used opioids, and some people chose not to respond. Does the fact that they did not respond provide extra information? Should we treat them equivalently as never-users?

This approach essentially creates a third group for this predictor: the “did not respond” group.

Types of Missingness

Sources of Missingness

Missing data can arise from various places in data:

- A survey was conducted and values were just randomly missed when being entered in the computer.
- A respondent chooses not to respond to a question like 'Have you ever recreationally used opioids?'.
- You decide to start collecting a new variable (due to new actions: like a pandemic) partway through the data collection of a study.
- You want to measure the speed of meteors, and some observations are just 'too quick' to be measured properly.

The source of missing values in data can lead to the major types of missingness:

Types of Missingness

There are 3 major types of missingness to be concerned about:

1. **Missing Completely at Random (MCAR)** - the probability of missingness in a variable is the same for all units. Like randomly poking holes in a data set.
2. **Missing at Random (MAR)** - the probability of missingness in a variable depends only on available information (in other predictors).
3. **Missing Not at Random (MNAR)** - the probability of missingness depends on information that has not been recorded and this information also predicts the missing values.

What are examples of each these 3 types?

Missing completely at random (MCAR)

Missing Completely at Random is the best case scenario, and the easiest to handle:

- Examples: a coin is flipped to determine whether an entry is removed. Or when values were just randomly missed when being entered in the computer.
- Effect if you ignore: there is no effect on inferences (estimates of β).
- How to handle: lots of options, but best to impute (more on next slide).

Missing at random (MAR)

Missing at Random is still a case that can be handled.

- Example(s): men and women respond to the question "have you ever felt harassed at work?" at different rates (and may be harassed at different rates).
- Effect if you ignore: inferences are biased (estimates of β 's) and predictions are usually worsened.
- How to handle: use the information in the other predictors to build a model and **impute** a value for the missing entry.

Key: we can fix any biases by modeling and imputing the missing values based on what is observed!

Missing Not at Random (MNAR)

Missing Not at Random is the worst case scenario, and impossible to handle properly:

- Example(s): patients drop out of a study because they experience some really bad side effect that was not measured. Or cheaters are less likely to respond when asked if you've ever cheated.
- Effect if you ignore: there are major effects on inferences (estimates of beta) or predictions.
- How to handle: you can 'improve' things by dealing with it like it is MAR, but you [likely] may never completely fix the bias. And incorporating a **missingness indicator variable** may actually be the best approach (if it is in a predictor).

What type of missingness is present?

Can you ever tell based on your data what type of missingness is actually present?

Since we asked the question, the answer must be no.

It generally cannot be determined whether data really are missing at random, or whether the missingness depends on unobserved predictors or the missing data themselves. The problem is that these potential “lurking variables” are unobserved (by definition) and so can never be completely ruled out.

In practice, a model with as many predictors as possible is used so that the ‘missing at random’ assumption is reasonable.

Imputation Methods

Handling Missing Data

When encountering missing data, the approach to handling it depends on:

1. whether the missing values are in the response or in the predictors. Generally speaking, it is much easier to handle missingness in predictors.
2. whether the variable is quantitative or categorical.
3. how much missingness is present in the variable. If there is too much missingness, you may be doing more damage than good.

Generally speaking, it is a good idea to attempt to **impute** (or ‘fill in’) entries for missing values in a variable (assuming your method of imputation is a good one).

Imputation Methods








There are several different approaches to imputing missing values:

1. **Impute the mean or median** (quantitative) or most common class (categorical) for all missing values in a variable.
2. Create a new variable that is an **indicator of missingness**, and include it in any model to predict the response (also plug in zero or the mean in the actual variable).
3. **Hot deck imputation**: for each missing entry, randomly select an observed entry in the variable and plug it in.
4. **Model the imputation**: plug in predicted values (\hat{y}) from a model based on the other observed predictors.
5. **Model the imputation with uncertainty**: plug in predicted values plus randomness ($\hat{y} + \epsilon$) from a model based on the other observed predictors.

What are the advantages and disadvantages of each approach?








Schematic: imputation through modeling

How do we use models to fill in missing data?

| X | Y |
|---|------|
|  | 1 |
|  | ? |
|  | 0.5 |
|  | 0.1 |
|  | ? |
|  | 10 |
|  | 0.03 |


Schematic: imputation through modeling

How do we use models to fill in missing data?

| <u>X_train</u> | <u>Y_train</u> | <u>X_test</u> | <u>Y_pred</u> |
|---|----------------|--|---------------|
|  | 1 |  | ? |
|  | 0.5 | | |
|  | 0.1 | | |
|  | 10 |  | ? |
|  | 0.03 | | |








Schematic: imputation through modeling

How do we use models to fill in missing data? Using k -NN for $k = 2$?









Schematic: imputation through modeling

How do we use models to fill in missing data? Using k -NN for $k = 2$?

| X | Y |
|--|--------------------|
|  | 1 |
|  | ? = (1 + 0.5) / 2 |
|  | 0.5 |
|  | 0.1 |
|  | ? = (0.1 + 10) / 2 |
|  | 10 |
|  | 0.03 |

Schematic: imputation through modeling

How do we use models to fill in missing data? Using linear regression?

| X | Y |
|---|--|
|  | 1 |
|  | ? = $m \cdot \text{red circle} + b$ |
|  | 0.5 |
|  | 0.1 |
|  | ? = $m \cdot \text{yellow circle} + b$ |
|  | 10 |
|  | 0.03 |

Where m and b are computed from the observations (rows) that do not have missingness (we should call them $b = \beta_0$ and $m = \beta_1$).


Imputation through modeling with uncertainty

The schematic in the last few slides ignores the fact of imputing with uncertainty. What happens if you ignore this fact and just use the 'best' model to impute values solely on \hat{y} ?

The distribution of the imputed values will be too narrow and not represent real data (see next slide for illustration). The goal is to impute values that include the uncertainty of the model.

How can this be done in practice in k -NN? In linear regression? In logistic regression?

Imputation through modeling with uncertainty: an illustration



Imputation through modeling with uncertainty: an illustration

Recall the probabilistic model in linear regression:

$$Y = \beta_0 + \beta_1 X_1 + \dots + \beta_p X_p + \varepsilon$$

where $\varepsilon \sim N(0, \sigma^2)$. How can we take advantage of this model to impute with uncertainty?

It's a 3 step process:

1. Fit a model to predict the predictor variable with missingness from all the other predictors.
2. Predict the missing values from the model in the previous part.
3. Add in a measure of uncertainty to this prediction by randomly sampling from a $N(0, \sigma^2)$ distribution, where σ^2 is the mean square error (MSE) from the model.

Imputation through modeling with uncertainty: k -NN regression

How can we use k -NN regression to impute values that mimic the error in our observations?

Two ways:

- Use $k = 1$.
- Use any other k , but randomly select from the nearest neighbors in \mathcal{N}_0 . This can be done with equal probability or with some weighting (inverse to the distance measure used).

Imputation through modeling with uncertainty: classifiers

For classifiers, this imputation with uncertainty/randomness is a little easier process. How can it be implemented?

If a classification model (logistic, k -NN, etc...) is used to predict the variable with missingness on the observed predictors, then all you need to do is flip a 'biased coin' (or multi-sided die) with the probabilities of coming up for each class equal to the predicted probabilities from the model.

Warning: do not just classify blindly using the predict command in sklearn!

Imputation across multiple variables

If only one variable has missing entries, life is easy. But what if all the predictor variables have a little bit of missingness (with some observations having multiple entries missing)? How can we handle that?

It's an iterative process. Impute X_1 based on X_2, \dots, X_p . Then impute X_2 based on X_1 and X_3, \dots, X_p . And continue down the line.

Any issues? Yes, not all of the missing values may be imputed with just one 'run' through the data set. So you will have to repeat these 'runs' until you have a completely filled in data set.

Multiple imputation: beyond this class

What is an issue with treating your now ‘complete’ data set (a mixture of actually observed values and imputed values) as simply all observed values?

Any inferences or predictions carried out will be tuned and potentially overfit to the random entries imputed for the missing entries. How can we prevent this phenomenon?

By performing **multiple imputation**: rerun the imputation algorithm many times, refit the model on the response many times (one time each), and then ‘average’ the predictions or estimates of β coefficients to perform inferences (also incorporating the uncertainty involved).


Note: this is beyond what we would expect in this class. But it generally a good thing to be aware of.

sklearn's impute module

Want to do imputation in sklearn? Of course there are functions for that!

The 4 main functions in this module are:

1. `sklearn.impute.SimpleImputer`: perform mean, median, mode, or any specific value to impute
2. `sklearn.impute.IterativeImputer`: perform multivariate imputation that estimates each predictor from all the others (user-supplied model/estimator).
3. `sklearn.impute.KNNImputer`: perform imputation automatically from a k -NN model from all the other predictors.
4. `sklearn.impute.MissingIndicator`: to create binary indicator(s) for where the missing values lie.



Exercise Time!

Ex. 1 (graded): Basic imputation (15+ min)

Ex. 2 (not graded): More complex imputation (30+ min)