

proj1

October 24, 2018

Before you turn this problem in, make sure everything runs as expected. First, **restart the kernel** (in the menubar, select Kernel→Restart) and then **run all cells** (in the menubar, select Cell→Run All).

Make sure you fill in any place that says YOUR CODE HERE or "YOUR ANSWER HERE", as well as your name and collaborators below:

```
In [1]: NAME = "Devin Hua"
        COLLABORATORS = ""
```

1 Project 1: Trump, Twitter, and Text

Welcome to the first project of Data 100! In this project, we will work with the Twitter API in order to analyze Donald Trump's tweets.

The project is due 11:59pm Tuesday, October 23, California Time.

You do not have to work on this project before the midterm, but you might find it helpful, since it goes over a lot of pandas materials that we haven't used in a while.

Fun:

We intend this project to be fun! You will analyze actual data from the Twitter API. You will also draw conclusions about the current (and often controversial) US President's tweet behavior. If you find yourself getting frustrated or stuck on one problem for too long, we suggest coming into office hours and working with friends in the class.

With that in mind, let's get started!

```
In [2]: # Run this cell to set up your notebook
import csv
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
import zipfile

# Ensure that Pandas shows at least 280 characters in columns, so we can see full tweets
pd.set_option('max_colwidth', 280)

%matplotlib inline
plt.style.use('fivethirtyeight')
```

```
import seaborn as sns
sns.set()
sns.set_context("talk")
import re
```

2 Downloading Recent Tweets

Since we'll be looking at Twitter data, we need to download the data from Twitter!

Twitter provides an API for downloading tweet data in large batches. The `tweepy` package makes it fairly easy to use.

```
In [3]: ## Make sure you are in your data100 conda environment if you are working locally.
        # The following should run:
import tweepy
```

There are instructions on using `tweepy` [here](#), but we will give you example code.

Twitter requires you to have authentication keys to access their API. To get your keys, you'll have to sign up as a Twitter developer. The next question will walk you through this process.

2.1 Question 1

Follow the instructions below to get your Twitter API keys. **Read the instructions completely before starting.**

1. [Create a Twitter account](#). You can use an existing account if you have one; if you prefer to not do this assignment under your regular account, feel free to create a throw-away account.
2. Under account settings, add your phone number to the account.
3. [Create a Twitter developer account](#) by clicking the 'Apply' button on the top right of the page. Attach it to your Twitter account. You'll have to fill out a form describing what you want to do with the developer account. Explain that you are doing this for a class at UC Berkeley and that you don't know exactly what you're building yet and just need the account to get started. These applications are approved by some sort of AI system, so it doesn't matter exactly what you write. Just don't enter a bunch of alweiofalwihflawiuehflawuihflaiwhfe type stuff or you might get rejected.
4. Once you're logged into your developer account, [create an application for this assignment](#). You can call it whatever you want, and you can write any URL when it asks for a web site. You don't need to provide a callback URL.
5. On the page for that application, find your Consumer Key and Consumer Secret.
6. On the same page, create an Access Token. Record the resulting Access Token and Access Token Secret.
7. Edit the file `keys.json` and replace the placeholders with your keys.

2.2 WARNING (Please Read) !!!!

2.2.1 Protect your Twitter Keys

If someone has your authentication keys, they can access your Twitter account and post as you! So don't give them to anyone, and **don't write them down in this notebook**. The usual way to

store sensitive information like this is to put it in a separate file and read it programmatically. That way, you can share the rest of your code without sharing your keys. That's why we're asking you to put your keys in `keys.json` for this assignment.

2.2.2 Avoid making too many API calls.

Twitter limits developers to a certain rate of requests for data. If you make too many requests in a short period of time, you'll have to wait awhile (around 15 minutes) before you can make more. So carefully follow the code examples you see and don't rerun cells without thinking. Instead, always save the data you've collected to a file. We've provided templates to help you do that.

2.2.3 Be careful about which functions you call!

This API can retweet tweets, follow and unfollow people, and modify your twitter settings. Be careful which functions you invoke! One of the sp18 instructors accidentally re-tweeted some tweets because that instructor typed `retweet` instead of `retweet_count`.

```
In [4]: import json
        key_file = 'keys.json'
        # Loading your keys from keys.json (which you should have filled
        # in in question 1):
        with open(key_file) as f:
            keys = json.load(f)
        # if you print or view the contents of keys be sure to delete the cell!
```

This cell tests the Twitter authentication. It should run without errors or warnings and display your Twitter username.

```
In [5]: import tweepy
        from tweepy import TweepError
        import logging

        try:
            auth = tweepy.OAuthHandler(keys["consumer_key"], keys["consumer_secret"])
            auth.set_access_token(keys["access_token"], keys["access_token_secret"])
            api = tweepy.API(auth)
            print("Your username is:", api.auth.get_username())
        except TweepError as e:
            logging.warning("There was a Tweepy error. Double check your API keys and try again.")
            logging.warning(e)
```

Your username is: devinhua123

2.3 Question 2

In the example below, we have loaded some tweets by @BerkeleyData. Run it and read the code.

```

In [6]: from pathlib import Path
import json

ds_tweets_save_path = "BerkeleyData_recent_tweets.json"
# Guarding against attempts to download the data multiple
# times:
if not Path(ds_tweets_save_path).is_file():
    # Getting as many recent tweets by @BerkeleyData as Twitter will let us have.
    # We use tweet_mode='extended' so that Twitter gives us full 280 character tweets.
    # This was a change introduced in September 2017.

    # The tweepy Cursor API actually returns "sophisticated" Status objects but we
    # will use the basic Python dictionaries stored in the _json field.
    example_tweets = [t._json for t in tweepy.Cursor(api.user_timeline, id="BerkeleyData",
                                                    tweet_mode='extended').items()]

    # Saving the tweets to a json file on disk for future analysis
    with open(ds_tweets_save_path, "w") as f:
        json.dump(example_tweets, f)

    # Re-loading the json file:
    with open(ds_tweets_save_path, "r") as f:
        example_tweets = json.load(f)

```

Assuming everything ran correctly you should be able to look at the first tweet by running the cell below.

Warning Do not attempt to view all the tweets in a notebook. It will likely freeze your browser. The following would be a **bad idea**:

```
pprint(example_tweets)
```

```

In [7]: # Looking at one tweet object, which has type Status:
from pprint import pprint # ...to get a more easily-readable view.
pprint(example_tweets[0])

```

```

{'contributors': None,
 'coordinates': None,
 'created_at': 'Tue Oct 09 17:35:11 +0000 2018',
 'display_text_range': [0, 155],
 'entities': {'hashtags': [{'indices': [142, 155], 'text': 'BerkeleyMIDS'}],
              'media': [{'display_url': 'pic.twitter.com/U34sa3XYc7',
                          'expanded_url': 'https://twitter.com/BerkeleyData/status/1049714909011951616/photo/1',
                          'id': 1049714859665813504,
                          'id_str': '1049714859665813504',
                          'indices': [156, 179],
                          'media_url': 'http://pbs.twimg.com/media/DpFWGm3UcAAx_Jq.jpg',
                          'media_url_https': 'https://pbs.twimg.com/media/DpFWGm3UcAAx_Jq.jpg',
                          'sizes': {'large': {'h': 400,
                                              'resize': 'fit',

```

```

        'w': 720},
        'medium': {'h': 400,
                    'resize': 'fit',
                    'w': 720},
        'small': {'h': 378,
                  'resize': 'fit',
                  'w': 680},
        'thumb': {'h': 150,
                  'resize': 'crop',
                  'w': 150}},
        'type': 'photo',
        'url': 'https://t.co/U34sa3XYc7'}],
'symbols': [],
'urls': [{'display_url': 'bit.ly/2MYD8l3',
          'expanded_url': 'https://bit.ly/2MYD8l3',
          'indices': [118, 141],
          'url': 'https://t.co/lNvOA36jqb'}],
'user_mentions': [],
'extended_entities': {'media': [{'display_url': 'pic.twitter.com/U34sa3XYc7',
                                  'expanded_url': 'https://twitter.com/BerkeleyData/status/1049714909011951616/photo',
                                  'id': 1049714859665813504,
                                  'id_str': '1049714859665813504',
                                  'indices': [156, 179],
                                  'media_url': 'http://pbs.twimg.com/media/DpFWGm3UcAAx_Jq.jpg',
                                  'media_url_https': 'https://pbs.twimg.com/media/DpFWGm3UcAAx_Jq.jpg',
                                  'sizes': {'large': {'h': 400,
                                                       'resize': 'fit',
                                                       'w': 720},
                                             'medium': {'h': 400,
                                                         'resize': 'fit',
                                                         'w': 720},
                                             'small': {'h': 378,
                                                         'resize': 'fit',
                                                         'w': 680},
                                             'thumb': {'h': 150,
                                                         'resize': 'crop',
                                                         'w': 150}}},
                                  'type': 'photo',
                                  'url': 'https://t.co/U34sa3XYc7'}]},
'favorite_count': 3,
'favorited': False,
'full_text': 'Meet Ramya Balasubramaniam, a datascience@berkeley student, and '
            'Control and Automation Engineer from Toronto, Canada: '
            'https://t.co/lNvOA36jqb #BerkeleyMIDS https://t.co/U34sa3XYc7',
'geo': None,
'id': 1049714909011951616,
'id_str': '1049714909011951616',
'in_reply_to_screen_name': None,

```

```

'in_reply_to_status_id': None,
'in_reply_to_status_id_str': None,
'in_reply_to_user_id': None,
'in_reply_to_user_id_str': None,
'is_quote_status': False,
'lang': 'en',
'place': None,
'possibly_sensitive': False,
'retweet_count': 2,
'retweeted': False,
'source': '<a href="http://twitter.com" rel="nofollow">Twitter Web Client</a>',
'truncated': False,
'user': {'contributors_enabled': False,
        'created_at': 'Thu Feb 28 14:37:26 +0000 2013',
        'default_profile': False,
        'default_profile_image': False,
        'description': 'An online Master of Information and Data Science '
                        '(MIDS) degree from the UC Berkeley School of '
                        'Information. Learn more at: http://t.co/zf6gfBWovQ',
        'entities': {'description': {'urls': [{'display_url': 'bit.ly/tBerkeleyData',
                                                'expanded_url': 'http://bit.ly/tBerkeleyData',
                                                'indices': [122, 144],
                                                'url': 'http://t.co/zf6gfBWovQ'}]}},
                    'url': {'urls': [{'display_url': 'datascience.berkeley.edu',
                                                'expanded_url': 'http://datascience.berkeley.edu',
                                                'indices': [0, 22],
                                                'url': 'http://t.co/S79Ul3oCaa'}]}},
        'favourites_count': 166,
        'follow_request_sent': False,
        'followers_count': 11739,
        'following': False,
        'friends_count': 412,
        'geo_enabled': False,
        'has_extended_profile': False,
        'id': 1227698863,
        'id_str': '1227698863',
        'is_translation_enabled': False,
        'is_translator': False,
        'lang': 'en',
        'listed_count': 485,
        'location': 'Berkeley, CA',
        'name': 'datascience@berkeley',
        'notifications': False,
        'profile_background_color': 'CCCCCC',
        'profile_background_image_url': 'http://abs.twimg.com/images/themes/theme1/bg.png',
        'profile_background_image_url_https': 'https://abs.twimg.com/images/themes/theme1/bg.png',
        'profile_background_tile': False,
        'profile_banner_url': 'https://pbs.twimg.com/profile_banners/1227698863/1502212054',

```

```

'profile_image_url': 'http://pbs.twimg.com/profile_images/894968224973897728/II8iiF3J_normal.jpg',
'profile_image_url_https': 'https://pbs.twimg.com/profile_images/894968224973897728/II8iiF3J_normal.jpg',
'profile_link_color': '5173B6',
'profile_sidebar_border_color': 'FFFFFF',
'profile_sidebar_fill_color': 'DDEEF6',
'profile_text_color': '333333',
'profile_use_background_image': True,
'protected': False,
'screen_name': 'BerkeleyData',
'statuses_count': 2403,
'time_zone': None,
'translator_type': 'none',
'url': 'http://t.co/S79U13oCaa',
'utc_offset': None,
'verified': False}}

```

2.4 Question 2a

2.4.1 What you need to do.

Re-factor the above code fragment into reusable snippets below. You should not need to make major modifications; this is mostly an exercise in understanding the above code block.

```

In [8]: def load_keys(path):
        """Loads your Twitter authentication keys from a file on disk.

        Args:
            path (str): The path to your key file. The file should
                be in JSON format and look like this (but filled in):
                {
                    "consumer_key": "<your Consumer Key here>",
                    "consumer_secret": "<your Consumer Secret here>",
                    "access_token": "<your Access Token here>",
                    "access_token_secret": "<your Access Token Secret here>"
                }

        Returns:
            dict: A dictionary mapping key names (like "consumer_key") to
                key values."""

        # YOUR CODE HERE
        key_file = 'keys.json'
        with open(key_file) as f:
            keys = json.load(f)
        return keys
        #raise NotImplementedError()

```

```

In [9]: def download_recent_tweets_by_user(user_account_name, keys):

```

```
"""Downloads tweets by one Twitter user.
```

Args:

```
user_account_name (str): The name of the Twitter account
    whose tweets will be downloaded.
keys (dict): A Python dictionary with Twitter authentication
    keys (strings), like this (but filled in):
    {
        "consumer_key": "<your Consumer Key here>",
        "consumer_secret": "<your Consumer Secret here>",
        "access_token": "<your Access Token here>",
        "access_token_secret": "<your Access Token Secret here>"
    }
```

Returns:

```
list: A list of Dictionary objects, each representing one tweet."""
```

```
import tweepy
auth = tweepy.OAuthHandler(keys["consumer_key"], keys["consumer_secret"])
auth.set_access_token(keys["access_token"], keys["access_token_secret"])
api = tweepy.API(auth)
tweets = [t._json for t in tweepy.Cursor(api.user_timeline, id=user_account_name,
                                         tweet_mode='extended').items()]

return tweets
# YOUR CODE HERE
#raise NotImplementedError()
```

In [10]: `def save_tweets(tweets, path):`

```
    """Saves a list of tweets to a file in the local filesystem.
```

This function makes no guarantee about the format of the saved tweets, ****except**** that calling `load_tweets(path)` after `save_tweets(tweets, path)` will produce the same list of tweets and that only the file at the given path is used to store the tweets. (That means you can implement this function however you want, as long as saving and loading works!)

Args:

```
tweets (list): A list of tweet objects (of type Dictionary) to
    be saved.
path (str): The place where the tweets will be saved.
```

Returns:

```
None"""
# YOUR CODE HERE

with open(path, "w") as f:
    json.dump(tweets, f)
#raise NotImplementedError()
```



```
In [11]: def load_tweets(path):
        """Loads tweets that have previously been saved.

        Calling load_tweets(path) after save_tweets(tweets, path)
        will produce the same list of tweets.

        Args:
            path (str): The place where the tweets were be saved.

        Returns:
            list: A list of Dictionary objects, each representing one tweet."""

        # YOUR CODE HERE
        with open(path, "r") as f:
            tweets = json.load(f)
        return tweets
        #raise NotImplementedError()
```

```
In [12]: def get_tweets_with_cache(user_account_name, keys_path):
        """Get recent tweets from one user, loading from a disk cache if available.
```

The first time you call this function, it will download tweets by a user. Subsequent calls will not re-download the tweets; instead they'll load the tweets from a save file in your local filesystem. All this is done using the functions you defined in the previous cell. This has benefits and drawbacks that often appear when you cache data:

- + : Using this function will prevent extraneous usage of the Twitter API.
- + : You will get your data much faster after the first time it's called.
- : If you really want to re-download the tweets (say, to get newer ones, or because you screwed up something in the previous cell and your tweets aren't what you wanted), you'll have to find the save file (which will look like <something>_recent_tweets.pkl) and delete it.

```
Args:
    user_account_name (str): The Twitter handle of a user, without the @.
    keys_path (str): The path to a JSON keys file in your filesystem.
"""
```

```
# YOUR CODE HERE
keys_data = load_keys(keys_path)
save_path = user_account_name + "_recent_tweets.json"
if not Path(save_path).is_file():
    downloaded_data = download_recent_tweets_by_user(user_account_name, keys_data)
    save_tweets(downloaded_data, save_path)
    return downloaded_data
else:
    return load_tweets(save_path)
```

```
#raise NotImplementedError()
```

If everything was implemented correctly you should be able to obtain roughly the last 3000 tweets by the realdonaldtrump. (This may take a few minutes)

```
In [13]: # When you are done, run this cell to load @realdonaldtrump's tweets.
# Note the function get_tweets_with_cache. You may find it useful
# later.
trump_tweets = get_tweets_with_cache("realdonaldtrump", key_file)
print("Number of tweets downloaded:", len(trump_tweets))
```

Number of tweets downloaded: 3220

```
In [14]: assert 2000 <= len(trump_tweets) <= 4000
```

2.4.2 Question 2b

We are limited to how many tweets we can download. In what month is the oldest tweet from Trump?

```
In [15]: # Enter the number of the month of the oldest tweet (e.g. 1 for January)
oldest_month = 10

# YOUR CODE HERE
#raise NotImplementedError()
```

2.5 Question 3

IMPORTANT! PLEASE READ

Unfortunately, Twitter prevent us from going further back in time using the public APIs. Fortunately, we have a snapshot of earlier tweets that we can combine with our new data.

We will again use the `fetch_and_cache` utility to download the dataset.

```
In [16]: # Download the dataset
from utils import fetch_and_cache
data_url = 'http://www.ds100.org/fa18/assets/datasets/old_trump_tweets.json.zip'
file_name = 'old_trump_tweets.json.zip'

dest_path = fetch_and_cache(data_url=data_url, file=file_name)
print(f'Located at {dest_path}')
```

Using version already downloaded: Fri Oct 12 21:17:12 2018

MD5 hash of file: b6e33874de91d1a40207cdf9f9b51a09

Located at data/old_trump_tweets.json.zip

Finally, we we will load the tweets directly from the compressed file without decompressing it first.

```
In [17]: my_zip = zipfile.ZipFile(dest_path, 'r')
         with my_zip.open("old_trump_tweets.json", "r") as f:
             old_trump_tweets = json.load(f)
```

This data is formatted identically to the recent tweets we just downloaded:

```
In [18]: pprint(old_trump_tweets[0])
```

```
{'contributors': None,
 'coordinates': None,
 'created_at': 'Wed Oct 12 14:00:48 +0000 2016',
 'entities': {'hashtags': [{'indices': [23, 38], 'text': 'CrookedHillary'}],
               'media': [{'display_url': 'pic.twitter.com/wjsl8ITVvk',
                           'expanded_url': 'https://twitter.com/realDonaldTrump/status/786204978629185536/video/1',
                           'id': 786204885318561792,
                           'id_str': '786204885318561792',
                           'indices': [39, 62],
                           'media_url': 'http://pbs.twimg.com/ext_tw_video_thumb/786204885318561792/pu/img/XqM...',
                           'media_url_https': 'https://pbs.twimg.com/ext_tw_video_thumb/786204885318561792/pu/i...',
                           'sizes': {'large': {'h': 576,
                                                'resize': 'fit',
                                                'w': 1024},
                                     'medium': {'h': 338,
                                                'resize': 'fit',
                                                'w': 600},
                                     'small': {'h': 191,
                                                'resize': 'fit',
                                                'w': 340},
                                     'thumb': {'h': 150,
                                                'resize': 'crop',
                                                'w': 150}}},
                           'type': 'photo',
                           'url': 'https://t.co/wjsl8ITVvk'}],
               'symbols': [],
               'urls': [],
               'user_mentions': []},
 'extended_entities': {'media': [{'additional_media_info': {'monetizable': False},
                                'display_url': 'pic.twitter.com/wjsl8ITVvk',
                                'expanded_url': 'https://twitter.com/realDonaldTrump/status/786204978629185536/vi...',
                                'id': 786204885318561792,
                                'id_str': '786204885318561792',
                                'indices': [39, 62],
                                'media_url': 'http://pbs.twimg.com/ext_tw_video_thumb/786204885318561792/pu/in...',
                                'media_url_https': 'https://pbs.twimg.com/ext_tw_video_thumb/78620488531856179...',
                                'sizes': {'large': {'h': 576,
                                                    'resize': 'fit',
                                                    'w': 1024},
                                          'medium': {'h': 338,
```

```

        'resize': 'fit',
        'w': 600},
    'small': {'h': 191,
        'resize': 'fit',
        'w': 340},
    'thumb': {'h': 150,
        'resize': 'crop',
        'w': 150}},
    'type': 'video',
    'url': 'https://t.co/wjsl8ITVvk',
    'video_info': {'aspect_ratio': [16, 9],
        'duration_millis': 30106,
        'variants': [{'bitrate': 832000,
            'content_type': 'video/mp4',
            'url': 'https://video.twimg.com/ext_tw_video/7862048853185617',
            {'bitrate': 2176000,
                'content_type': 'video/mp4',
                'url': 'https://video.twimg.com/ext_tw_video/7862048853185617',
                {'bitrate': 320000,
                    'content_type': 'video/mp4',
                    'url': 'https://video.twimg.com/ext_tw_video/7862048853185617',
                    {'content_type': 'application/x-mpegURL',
                        'url': 'https://video.twimg.com/ext_tw_video/7862048853185617'}]}]}]}],
    'favorite_count': 42242,
    'favorited': False,
    'geo': None,
    'id': 786204978629185536,
    'id_str': '786204978629185536',
    'in_reply_to_screen_name': None,
    'in_reply_to_status_id': None,
    'in_reply_to_status_id_str': None,
    'in_reply_to_user_id': None,
    'in_reply_to_user_id_str': None,
    'is_quote_status': False,
    'lang': 'en',
    'place': {'attributes': {},
        'bounding_box': {'coordinates': [[[-87.634643, 24.396308],
            [-79.974307, 24.396308],
            [-79.974307, 31.001056],
            [-87.634643, 31.001056]]],
            'type': 'Polygon'},
        'contained_within': [],
        'country': 'United States',
        'country_code': 'US',
        'full_name': 'Florida, USA',
        'id': '4ec01c9dbc693497',
        'name': 'Florida',
        'place_type': 'admin',

```

```

    'url': 'https://api.twitter.com/1.1/geo/id/4ec01c9dbc693497.json'},
'possibly_sensitive': False,
'retweet_count': 24915,
'retweeted': False,
'source': '<a href="http://twitter.com/download/iphone" '
    'rel="nofollow">Twitter for iPhone</a>',
'text': 'PAY TO PLAY POLITICS. \n#CrookedHillary https://t.co/wjsl8ITVvk',
'truncated': False,
'user': {'contributors_enabled': False,
    'created_at': 'Wed Mar 18 13:46:38 +0000 2009',
    'default_profile': False,
    'default_profile_image': False,
    'description': '45th President of the United States of America',
    'entities': {'description': {'urls': []}},
    'favourites_count': 12,
    'follow_request_sent': False,
    'followers_count': 35307313,
    'following': False,
    'friends_count': 45,
    'geo_enabled': True,
    'has_extended_profile': False,
    'id': 25073877,
    'id_str': '25073877',
    'is_translation_enabled': True,
    'is_translator': False,
    'lang': 'en',
    'listed_count': 74225,
    'location': 'Washington, DC',
    'name': 'Donald J. Trump',
    'notifications': False,
    'profile_background_color': '6D5C18',
    'profile_background_image_url': 'http://pbs.twimg.com/profile_background_images/530021613/trump_...',
    'profile_background_image_url_https': 'https://pbs.twimg.com/profile_background_images/530021613/trump_...',
    'profile_background_tile': True,
    'profile_banner_url': 'https://pbs.twimg.com/profile_banners/25073877/1501916634',
    'profile_image_url': 'http://pbs.twimg.com/profile_images/874276197357596672/kUuht00m_normal.jpg',
    'profile_image_url_https': 'https://pbs.twimg.com/profile_images/874276197357596672/kUuht00m_normal.jpg',
    'profile_link_color': '1B95E0',
    'profile_sidebar_border_color': 'BDDCAD',
    'profile_sidebar_fill_color': 'C5CEC0',
    'profile_text_color': '333333',
    'profile_use_background_image': True,
    'protected': False,
    'screen_name': 'realDonaldTrump',
    'statuses_count': 35480,
    'time_zone': 'Eastern Time (US & Canada)',
    'translator_type': 'regular',
    'url': None,

```

```
'utc_offset': -14400,
'verified': True}}
```

As a dictionary we can also list the keys:

```
In [19]: old_trump_tweets[0].keys()
```

Out[19]: dict_keys(['created_at', 'id', 'id_str', 'text', 'truncated', 'entities', 'extended_entities', 'source', 'in_reply_to_status_id_str'])

Since we're giving you a zipfile of old tweets, you may wonder why we didn't just give you a zipfile of ALL tweets and save you the trouble of creating a Twitter developer account. The reason is that we wanted you to see what it's like to collect data from the real world on your own. It can be a pain!

2.5.1 Question 3a

Merge the `old_trump_tweets` and the `trump_tweets` we downloaded from twitter into one giant list of tweets.

Important: There may be some overlap so be sure to eliminate duplicate tweets.

Hint: the id of a tweet is always unique.

```
In [20]: all_tweets = []
         old_trump_id = [ID['id'] for ID in old_trump_tweets]
         for elem in trump_tweets:
             if elem['id'] not in old_trump_id:
                 all_tweets.append(elem)
         all_tweets = all_tweets + old_trump_tweets
         # YOUR CODE HERE
         #raise NotImplementedError()
```

```
In [21]: assert len(all_tweets) > len(trump_tweets)
         assert len(all_tweets) > len(old_trump_tweets)
```

2.5.2 Question 3b

Construct a DataFrame called `trump` containing all the tweets stored in `all_tweets`. The index of the dataframe should be the ID of each tweet (looks something like 907698529606541312). It should have these columns:

- **time:** The time the tweet was created encoded as a datetime object. (Use `pd.to_datetime` to encode the timestamp.)
- **source:** The source device of the tweet.
- **text:** The text of the tweet.
- **retweet_count:** The retweet count of the tweet.

Finally, the resulting dataframe should be sorted by the index.

Warning: Some tweets will store the text in the `text` field and other will use the full `text` field.

```
In [22]: trump = pd.DataFrame(data=all_tweets, columns= ['created_at', 'source', 'text', 'full_text', 'retweet_count'])
trump['text'] = trump['text'].combine_first(trump['full_text'])
trump = trump.drop(columns='full_text')
trump['created_at'] = pd.to_datetime(trump['created_at'])
trump.columns = ['time', 'source', 'text', 'retweet_count', 'id']
trump = trump.set_index('id')
# YOUR CODE HERE
#raise NotImplementedError()

In [23]: assert isinstance(trump, pd.DataFrame)
assert trump.shape[0] < 11000
assert trump.shape[1] >= 4
assert 831846101179314177 in trump.index
assert 753063644578144260 in trump.index
assert all(col in trump.columns for col in ['time', 'source', 'text', 'retweet_count'])
# If you fail these tests, you probably tried to use __dict__ or _json to read in the tweets
assert np.sometrue([(('Twitter for iPhone' in s) for s in trump['source'].unique())])
assert trump['time'].dtype == np.dtype('<M8[ns]')
assert trump['text'].dtype == np.dtype('O')
assert trump['retweet_count'].dtype == np.dtype('int64')
```

2.6 Question 4: Tweet Source Analysis

In the following questions, we are going to find out the characteristics of Trump tweets and the devices used for the tweets.

First let's examine the source field:

```
In [24]: trump['source'].unique()

Out[24]: array(['<a href="http://twitter.com/download/iphone" rel="nofollow">Twitter for iPhone</a>',
                '<a href="https://studio.twitter.com" rel="nofollow">Media Studio</a>',
                '<a href="http://twitter.com/#!/download/ipad" rel="nofollow">Twitter for iPad</a>',
                '<a href="http://twitter.com" rel="nofollow">Twitter Web Client</a>',
                '<a href="http://twitter.com/download/android" rel="nofollow">Twitter for Android</a>',
                '<a href="http://instagram.com" rel="nofollow">Instagram</a>',
                '<a href="https://mobile.twitter.com" rel="nofollow">Mobile Web (M5)</a>',
                '<a href="https://ads.twitter.com" rel="nofollow">Twitter Ads</a>',
                '<a href="https://periscope.tv" rel="nofollow">Periscope</a>'], dtype=object)
```

2.7 Question 4a

Remove the HTML tags from the source field.

Hint: Use `trump['source'].str.replace` and your favorite regular expression.

```
In [25]: ## Uncomment and complete
trump['source'] = trump['source'].str.replace(r"<.*?>", "")


# YOUR CODE HERE
#raise NotImplementedError()
```

```
In [26]: from datetime import datetime
ELEC_DATE = datetime(2016, 11, 8)
INAUG_DATE = datetime(2017, 1, 20)
assert set(trump[(trump['time'] > ELEC_DATE) & (trump['time'] < INAUG_DATE)]['source'].unique()
           ['Twitter Web Client',
            'Twitter for Android',
            'Twitter for iPhone'])
```

We can see in the following plot that there are two device types that are more commonly used

```
In [27]: trump['source'].value_counts().plot(kind="bar")
plt.ylabel("Number of Tweets")
```

```
Out[27]: Text(0,0.5,'Number of Tweets')
```



2.8 Question 4b

Is there a difference between his Tweet behavior across these devices? We will attempt to answer this question in our subsequent analysis.

First, we'll take a look at whether Trump's tweets from an Android come at different times than his tweets from an iPhone. Note that Twitter gives us his tweets in the [UTC timezone](#) (notice the +0000 in the first few tweets)

```
In [28]: for t in trump_tweets[0:3]:
         print(t['created_at'])
```

```
Fri Oct 12 19:37:27 +0000 2018
Fri Oct 12 19:10:51 +0000 2018
Fri Oct 12 18:09:17 +0000 2018
```

We'll convert the tweet times to US Eastern Time, the timezone of New York and Washington D.C., since those are the places we would expect the most tweet activity from Trump.

```
In [29]: trump['est_time'] = (
         trump['time'].dt.tz_localize("UTC") # Set initial timezone to UTC
         .dt.tz_convert("EST") # Convert to Eastern Time
         )
         trump.head()
```

```
Out[29]:
```

	time	source \
id		
1050832839678353410	2018-10-12 19:37:27	Twitter for iPhone
1050826148131860480	2018-10-12 19:10:51	Twitter for iPhone
1050810652271529984	2018-10-12 18:09:17	Twitter for iPhone
1050788995377049601	2018-10-12 16:43:13	Twitter for iPhone
1050777580553588738	2018-10-12 15:57:52	Twitter for iPhone

```
id
1050832839678353410
1050826148131860480
1050810652271529984 People have no idea how hard Hurricane Michael has hit the great state of Georgia.
1050788995377049601
1050777580553588738
```

```
retweet_count    est_time
id
1050832839678353410    3155 2018-10-12 14:37:27-05:00
1050826148131860480    4531 2018-10-12 14:10:51-05:00
1050810652271529984    8375 2018-10-12 13:09:17-05:00
1050788995377049601    11004 2018-10-12 11:43:13-05:00
1050777580553588738    8797 2018-10-12 10:57:52-05:00
```

What you need to do:

Add a column called hour to the trump table which contains the hour of the day as floating point number computed by:

$$\text{hour} + \frac{\text{minute}}{60} + \frac{\text{second}}{60^2}$$

```
In [30]: trump['hour'] = trump['est_time'].dt.hour + trump['est_time'].dt.minute / 60 + trump['est_time'].dt.second / 3600
# make a bar plot here
# YOUR CODE HERE
#raise NotImplementedError()
```

```
In [31]: assert np.isclose(trump.loc[690171032150237184]['hour'], 8.93639)
```

2.9 Question 4c


Use this data along with the seaborn distplot function to examine the distribution over hours of the day in eastern time that trump tweets on each device for the 2 most commonly used devices. Your plot should look similar to the following.

```
In [32]: trump['source'].value_counts().head(2)
iphone = trump.loc[trump['source'] == 'Twitter for iPhone']
android = trump.loc[trump['source'] == 'Twitter for Android']

In [33]: ### make your plot here
# YOUR CODE HERE
sns.distplot(iphone['hour'], hist=False, label='iPhone')
sns.distplot(android['hour'], hist=False, label='Android')
plt.legend()
plt.ylabel('fraction')

#raise NotImplementedError()
```

```
Out[33]: Text(0,0.5,'fraction')
```



2.10 Question 4d

According to [this Verge article](#), Donald Trump switched from an Android to an iPhone sometime in March 2017.


Create a figure identical to your figure from 4c, except that you should show the results only from 2016.

During the campaign, it was theorized that Donald Trump's tweets from Android were written by him personally, and the tweets from iPhone were from his staff. Does your figure give support to this theory?

```
In [34]: iphone_2016 = iphone.loc[iphone['time'].dt.year == 2016]
        android_2016 = android.loc[android['time'].dt.year == 2016]
```

```
In [35]: ### make your plot here
sns.distplot(iphone_2016['hour'], hist=False, label='iPhone')
sns.distplot(android_2016['hour'], hist=False, label='Android')
plt.legend()
plt.ylabel('fraction')
# YOUR CODE HERE
#raise NotImplementedError()
```

```
Out[35]: Text(0,0.5,'fraction')
```



From the graph in 4c, the iPhone and Android line plots share very similar lines in the distplot. This indicates that the amount of tweets tweeted at their respective hours are the same for both iPhone and Android. However, before Trump's switch from Android to iPhone in 2017, in the year of 2016, represented by the plot in 4d, you can easily see the difference in the fraction of

tweets and the respective hours the tweets were tweeted. This implies that there's a third party tweeting on Trump's account for the iPhone since the Android line on the 2016 distplot remained the same.

2.11 Question 5

Let's now look at which device he has used over the entire time period of this dataset.

To examine the distribution of dates we will convert the date to a fractional year that can be plotted as a distribution.

(Code borrowed from <https://stackoverflow.com/questions/6451655/python-how-to-convert-datetime-dates-to-decimal-years>)

```
In [36]: import datetime
def year_fraction(date):
    start = datetime.date(date.year, 1, 1).toordinal()
    year_length = datetime.date(date.year+1, 1, 1).toordinal() - start
    return date.year + float(date.toordinal() - start) / year_length
```

```
trump['year'] = trump['time'].apply(year_fraction)
```


2.11.1 Question 5a

Use the `sns.distplot` to overlay the distributions of the 2 most frequently used web technologies over the years. Your final plot should look like:

```
In [37]: # new dataframe with year
iphone_with_year = trump.loc[trump['source'] == 'Twitter for iPhone']
android_with_year = trump.loc[trump['source'] == 'Twitter for Android']

# distplots
sns.distplot(iphone_with_year['year'], label='iphone')
sns.distplot(android_with_year['year'], label='android')
plt.legend()
#raise NotImplementedError()
```

```
Out[37]: <matplotlib.legend.Legend at 0x7f956f445ef0>
```



2.12 Question 6: Sentiment Analysis

It turns out that we can use the words in Trump's tweets to calculate a measure of the sentiment of the tweet. For example, the sentence "I love America!" has positive sentiment, whereas the sentence "I hate taxes!" has a negative sentiment. In addition, some words have stronger positive / negative sentiment than others: "I love America." is more positive than "I like America."

We will use the [VADER \(Valence Aware Dictionary and sEntiment Reasoner\)](#) lexicon to analyze the sentiment of Trump's tweets. VADER is a lexicon and rule-based sentiment analysis tool that is specifically attuned to sentiments expressed in social media which is great for our usage.

The VADER lexicon gives the sentiment of individual words. Run the following cell to show the first few rows of the lexicon:

```
In [38]: print(''.join(open("vader_lexicon.txt").readlines()[:10]))
```

```
$:      -1.5      0.80623      [-1, -1, -1, -1, -3, -1, -3, -1, -2, -1]
%)      -0.4      1.0198      [-1, 0, -1, 0, 0, -2, -1, 2, -1, 0]
%-)     -1.5      1.43178      [-2, 0, -2, -2, -1, 2, -2, -3, -2, -3]
&-:     -0.4      1.42829      [-3, -1, 0, 0, -1, -1, -1, 2, -1, 2]
&:     -0.7      0.64031      [0, -1, -1, -1, 1, -1, -1, -1, -1, -1]
( '{' )    1.6      0.66332      [1, 2, 2, 1, 1, 2, 2, 1, 3, 1]
( %      -0.9      0.9434      [0, 0, 1, -1, -1, -1, -2, -2, -1, -2]
(' -:     2.2      1.16619      [4, 1, 4, 3, 1, 2, 3, 1, 2, 1]
(' :     2.3      0.9        [1, 3, 3, 2, 2, 4, 2, 3, 1, 2]
(( -:     2.1      0.53852      [2, 2, 2, 1, 2, 3, 2, 2, 3, 2]
```

2.13 Question 6a

As you can see, the lexicon contains emojis too! The first column of the lexicon is the *token*, or the word itself. The second column is the *polarity* of the word, or how positive / negative it is.

(How did they decide the polarities of these words? What are the other two columns in the lexicon? See the link above.)

Read in the lexicon into a DataFrame called `sent`. The index of the DF should be the tokens in the lexicon. `sent` should have one column: `polarity`: The polarity of each token.

```
In [39]: sent = pd.read_csv('vader_lexicon.txt', sep=" ", header=None, delimiter="\t")
sent = sent.drop(columns=[2, 3])
sent = sent.set_index(0)
sent.columns = ['polarity']

# YOUR CODE HERE
#raise NotImplementedError()

In [40]: assert isinstance(sent, pd.DataFrame)
assert sent.shape == (7517, 1)
assert list(sent.index[5000:5005]) == ['paranoids', 'pardon', 'pardoned', 'pardoning', 'pardons']
assert np.allclose(sent['polarity'].head(), [-1.5, -0.4, -1.5, -0.4, -0.7])
```

2.14 Question 6b

Now, let's use this lexicon to calculate the overall sentiment for each of Trump's tweets. Here's the basic idea:

1. For each tweet, find the sentiment of each word.
2. Calculate the sentiment of each tweet by taking the sum of the sentiments of its words.

First, let's lowercase the text in the tweets since the lexicon is also lowercase. Set the `text` column of the `trump` DF to be the lowercased text of each tweet.

```
In [41]: # YOUR CODE HERE
trump['text'] = trump['text'].str.lower()
#raise NotImplementedError()

In [42]: assert trump['text'].loc[884740553040175104] == 'working hard to get the olympics for the united states (
```

2.15 Question 6c

Now, let's get rid of punctuation since it'll cause us to fail to match words. Create a new column called `no_punc` in the `trump` DF to be the lowercased text of each tweet with all punctuation replaced by a single space. We consider punctuation characters to be any character that isn't a Unicode word character or a whitespace character. You may want to consult the Python documentation on regexes for this problem.

(Why don't we simply remove punctuation instead of replacing with a space? See if you can figure this out by looking at the tweet data.)

```

In [43]: # Save your regex in punct_re
punct_re = r'^[\s\w]'
trump['no_punc'] = trump['text'].str.replace(punct_re, " ")

# YOUR CODE HERE
#raise NotImplementedError()

In [44]: assert isinstance(punct_re, str)
assert re.search(punct_re, 'this') is None
assert re.search(punct_re, 'this is ok') is None
assert re.search(punct_re, 'this is\nok') is None
assert re.search(punct_re, 'this is not ok.') is not None
assert re.search(punct_re, 'this#is#ok') is not None
assert re.search(punct_re, 'this^is ok') is not None
assert trump['no_punc'].loc[800329364986626048] == 'i watched parts of nbcnl saturday night live last n
assert trump['no_punc'].loc[894620077634592769] == 'on purpleheartday i thank all the brave men and w
# If you fail these tests, you accidentally changed the text column
assert trump['text'].loc[884740553040175104] == 'working hard to get the olympics for the united states (l

```

2.16 Question 6d:

Now, let's convert the tweets into what's called a *tidy format* to make the sentiments easier to calculate. Use the `no_punc` column of `trump` to create a table called `tidy_format`. The index of the table should be the IDs of the tweets, repeated once for every word in the tweet. It has two columns:

1. `num`: The location of the word in the tweet. For example, if the tweet was "i love america", then the location of the word "i" is 0, "love" is 1, and "america" is 2.
2. `word`: The individual words of each tweet.

The first few rows of our `tidy_format` table look like:

```

<tr style="text-align: right;">
  <th></th>
  <th>num</th>
  <th>word</th>
</tr>

<tr>
  <th>894661651760377856</th>
  <td>0</td>
  <td>i</td>
</tr>
<tr>
  <th>894661651760377856</th>
  <td>1</td>
  <td>think</td>
</tr>
<tr>

```

```

<th>894661651760377856</th>
<td>2</td>
<td>senator</td>
</tr>
<tr>
<th>894661651760377856</th>
<td>3</td>
<td>blumenthal</td>
</tr>
<tr>
<th>894661651760377856</th>
<td>4</td>
<td>should</td>
</tr>

```

Note that you'll get different results depending on when you pulled in the tweets. However, you can double check that your tweet with ID 894661651760377856 has the same rows as ours. Our tests don't check whether your table looks exactly like ours.

As usual, try to avoid using any for loops. Our solution uses a chain of 5 methods on the 'trump' DF, albeit using some rather advanced Pandas hacking.

- **Hint 1:** Try looking at the `expand` argument to `pandas' str.split`.
- **Hint 2:** Try looking at the `stack()` method.
- **Hint 3:** Try looking at the `level` parameter of the `reset_index` method.

```

In [45]: split = trump['no_punc'].str.split(expand=True)
         stack = split.stack()
         reset_index = stack.reset_index()
         tidy_format = reset_index.set_index('id')
         tidy_format.columns = ['num', 'word']

         # YOUR CODE HERE
         #raise NotImplementedError()

```

```

In [46]: assert tidy_format.loc[894661651760377856].shape == (27, 2)
         assert ' '.join(list(tidy_format.loc[894661651760377856]['word'])) == 'i think senator blumenthal should ta

```

2.17 Question 6e:

Now that we have this table in the tidy format, it becomes much easier to find the sentiment of each tweet: we can join the table with the lexicon table.

Add a polarity column to the trump table. The polarity column should contain the sum of the sentiment polarity of each word in the text of the tweet.

Hint you will need to merge the `tidy_format` and `sent` tables and group the final answer.

```

In [47]: merged = tidy_format.merge(sent, how='left', left_on='word', right_index = True)
         merged_sum = merged.groupby('id').agg(sum)

```



```

trump['polarity'] = merged_sum['polarity']

# YOUR CODE HERE
#raise NotImplementedError()

In [48]: assert np.allclose(trump.loc[744701872456536064, 'polarity'], 8.4)
         assert np.allclose(trump.loc[745304731346702336, 'polarity'], 2.5)
         assert np.allclose(trump.loc[744519497764184064, 'polarity'], 1.7)
         assert np.allclose(trump.loc[894661651760377856, 'polarity'], 0.2)
         assert np.allclose(trump.loc[894620077634592769, 'polarity'], 5.4)
         # If you fail this test, you dropped tweets with 0 polarity
         assert np.allclose(trump.loc[744355251365511169, 'polarity'], 0.0)

```

Now we have a measure of the sentiment of each of his tweets! Note that this calculation is rather basic; you can read over the VADER readme to understand a more robust sentiment analysis.

Now, run the cells below to see the most positive and most negative tweets from Trump in your dataset:

```

In [49]: print('Most negative tweets:')
         for t in trump.sort_values('polarity').head()['text']:
             print('\n ', t)

```

Most negative tweets:

it is outrageous that poisonous synthetic heroin fentanyl comes pouring into the u.s. postal system from china.
the rigged russian witch hunt goes on and on as the originators and founders of this scam continue to be fired a
james comey is a proven leaker & liar. virtually everyone in washington thought he should be fired for the
this is an illegally brought rigged witch hunt run by people who are totally corrupt and/or conflicted. it was sta
wheres the collusion? they made up a phony crime called collusion, and when there was no collusion they say th

```

In [50]: print('Most positive tweets:')
         for t in trump.sort_values('polarity', ascending=False).head()['text']:
             print('\n ', t)

```

Most positive tweets:

congratulations to patrick reed on his great and courageous masters win! when patrick had his amazing win at
my supporters are the smartest, strongest, most hard working and most loyal that we have seen in our countries
thank you to all of my great supporters, really big progress being made. other countries wanting to fix crazy tra
thank you, @wvgovernor jim justice, for that warm introduction. tonight, it was my great honor to attend the g

the republican party had a great night. tremendous voter energy and excitement, and all candidates are those v


2.18 Question 6g

Plot the distribution of tweet sentiments broken down by whether the text of the tweet contains nyt or fox. Then in the box below comment on what we observe?

```
In [51]: contains_nyt = trump.loc[trump['no_punc'].str.contains('nyt')]  
contains_fox = trump.loc[trump['no_punc'].str.contains('fox')]
```

```
In [52]: # YOUR CODE HERE  
sns.distplot(contains_nyt['polarity'], hist=False, label='nyt')  
sns.distplot(contains_fox['polarity'], hist=False, label='fox')  
plt.legend()  
#raise NotImplementedError()
```

```
Out[52]: <matplotlib.legend.Legend at 0x7f956e1406a0>
```



Comment on what you observe: Trump's tweets that include nyt are a lot more negative than his tweets that include fox. In general, the mean of the nyt tweets is around -5 and the mean for the fox tweets is around 2. The fox plot also has two peaks while the nyt plot has only one clear peak.

2.19 Question 7: Engagement

2.20 Question 7a

In this problem, we'll explore which words led to a greater average number of retweets. For example, at the time of this writing, Donald Trump has two tweets that contain the word 'oakland' (tweets 932570628451954688 and 1016609920031117312) with 36757 and 10286 retweets respectively, for an average of 23,521.5.

Find the top 20 most retweeted words. Include only words that appear in at least 25 tweets. As usual, try to do this without any for loops. You can string together ~7 pandas commands and get everything done on one line.

Your top_20 table should have this format:

```
<tr style="text-align: right;">
  <th></th>
  <th>retweet_count</th>
</tr>
<tr>
  <th>word</th>
  <th></th>
</tr>

<tr>
  <th>jong</th>
  <td>40675.666667</td>
</tr>
<tr>
  <th>try</th>
  <td>33937.800000</td>
</tr>
<tr>
  <th>kim</th>
  <td>32849.595745</td>
</tr>
<tr>
  <th>un</th>
  <td>32741.731707</td>
</tr>
<tr>
  <th>maybe</th>
  <td>30473.192308</td>
</tr>
```


```
In [53]: retweet_table = tidy_format.merge(trump[['retweet_count']], how='left', left_index=True, right_index =
retweet_table['occurrences'] = np.ones(len(retweet_table))
new_retweet_table = retweet_table.groupby('word').agg({'retweet_count': 'mean', 'occurrences': 'sum'})
filtered_table = new_retweet_table.loc[new_retweet_table['occurrences'] >= 25]
top_20 = filtered_table.sort_values('retweet_count', ascending=False).head(20).drop(columns='occurrences')
```

```
# YOUR CODE HERE
#raise NotImplementedError()
```

```
In [54]: # Although it can't be guaranteed, it's very likely that the top 5 words will still be
# in the top 20 words in the next month.
assert 'jong' in top_20.index
assert 'try' in top_20.index
assert 'kim' in top_20.index
assert 'un' in top_20.index
assert 'maybe' in top_20.index
```

Here's a bar chart of your results:

```
In [55]: top_20['retweet_count'].sort_values().plot.barh(figsize=(10, 8));
```



2.21 Question 7b

"kim", "jong" and "un" are apparently really popular in Trump's tweets! It seems like we can conclude that his tweets involving jong are more popular than his other tweets. Or can we?

Consider each of the statements about possible confounding factors below. State whether each statement is true or false and explain. If the statement is true, state whether the confounding factor could have made kim jong un related tweets higher in the list than they should be.

1. We didn't restrict our word list to nouns, so we have unhelpful words like "let" and "any" in our result.

2. We didn't remove hashtags in our text, so we have duplicate words (eg. #great and great).
 3. We didn't account for the fact that Trump's follower count has increased over time.
1. This statement is true, but the unhelpful words did not affect the kim jong un tweets by making them higher in the list than they should be. If anything, the unhelpful words should lower the kim jong un tweets due to the fact that these unhelpful words probably appear in a lot of the tweets, but in this case, it did not affect the related tweets at all and did make the tweets higher in the list than they should be.
 2. This statement is false; we removed hashtags in problem 6c.
 3. This statement is true and it is a confounding factor. In other words, this rise in followers could have made the kim jong un related tweets higher in the list than they should be. The counterpart is also true however. (The rise in followers could have also made other tweets rise up in the list)

2.22 Question 8

Using the trump tweets construct an interesting plot describing a property of the data and discuss what you found below.

Ideas:

1. How has the sentiment changed with length of the tweets?
2. Does sentiment affect retweet count?
3. Are retweets more negative than regular tweets?
4. Are there any spikes in the number of retweets and do they correspond to world events?
5. *Bonus:* How many Russian twitter bots follow Trump?
6. What terms have an especially positive or negative sentiment?

You can look at other data sources and even tweets.

2.22.1 Plot:

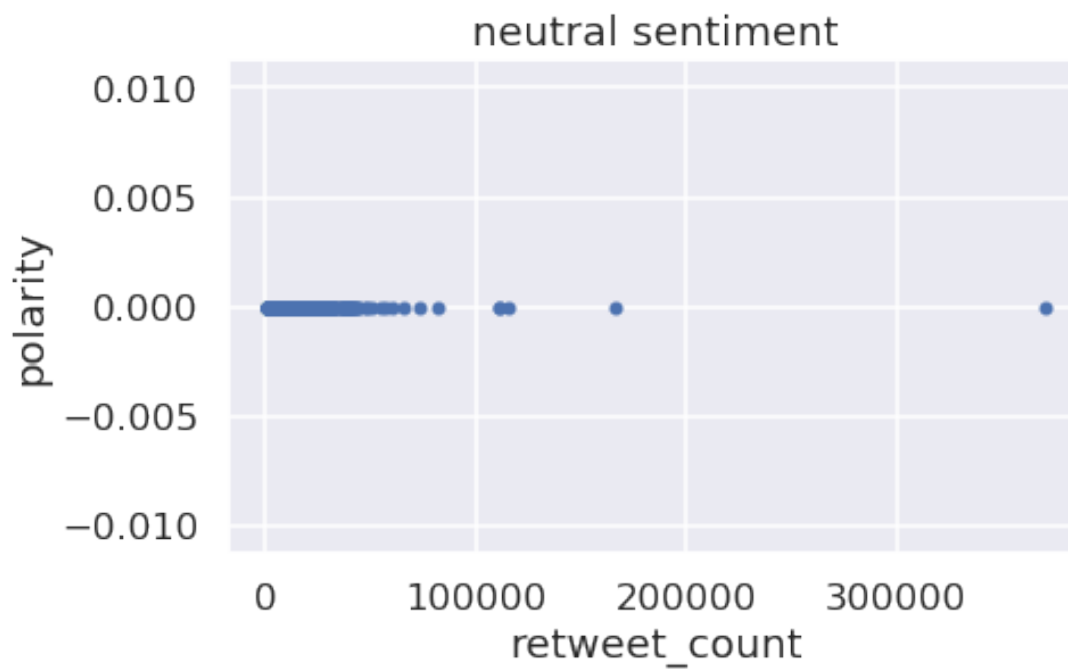

In [56]: # YOUR CODE HERE

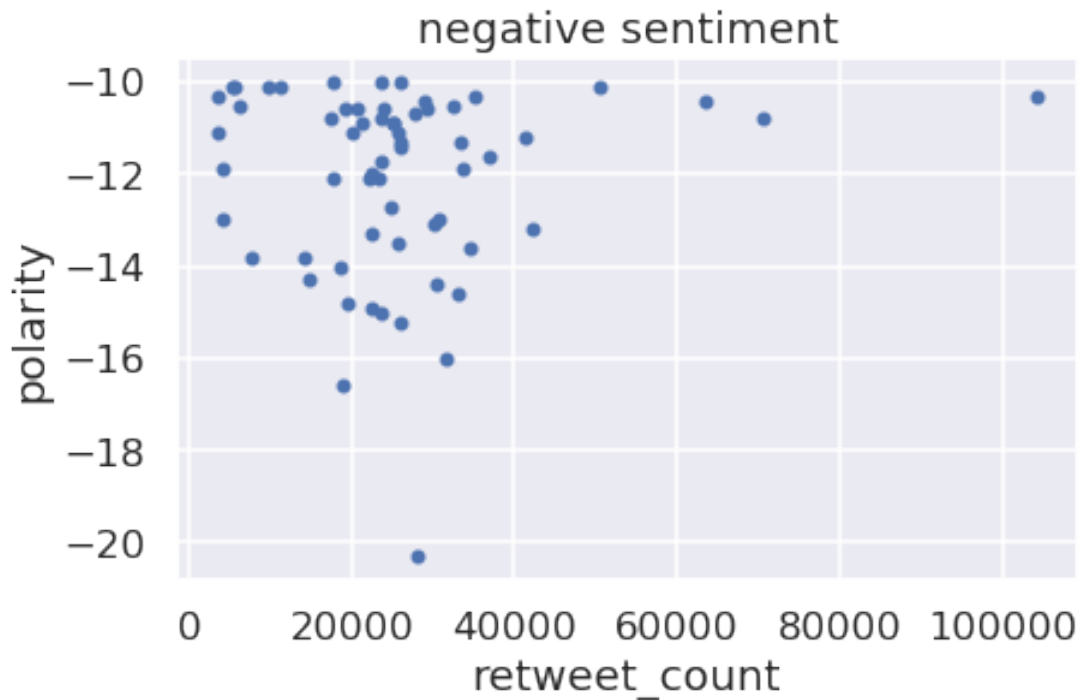
```
# new dataframes with sentiment value
sentiment_value = trump.filter(items=['retweet_count', 'polarity'])
positive_sent = sentiment_value.loc[sentiment_value['polarity'] >= 10]
neutral_sent = sentiment_value.loc[sentiment_value['polarity'] == 0]
negative_sent = sentiment_value.loc[sentiment_value['polarity'] <= -10]

#respective scatter plots
positive_sent.plot(x='retweet_count', y='polarity', kind='scatter', title='positive sentiment')
neutral_sent.plot(x='retweet_count', y='polarity', kind='scatter', title='neutral sentiment')
negative_sent.plot(x='retweet_count', y='polarity', kind='scatter', title='negative sentiment')

#raise NotImplementedError()
```

Out[56]: <matplotlib.axes._subplots.AxesSubplot at 0x7f956e839eb8>





2.22.2 Discussion of Your Plot:

I made three different scatter plots (positive, neutral, and negative) to see if sentiment affects the number of retweets. Just from looking at the scatter plots, there is no real significance between the sentiment value and the number of retweets. For all three scatter plots, it seems as if all the retweet counts lie at around 20,000 retweets with the positive and negative plots containing more outliers. This event could possibly be explained by the fact that Trump has a lot of followers on Twitter and that there would always be people who retweet his tweets just for the heck of it. One observation from the three scatter plots is that the neutral sentiment plot has fewer outliers than the ones with high/low polarity perhaps due to the fact that high/low polarity tweets contain words or phrases that are more profound or even take people by surprise. Another take from the scatter plots is that there are generally a large number of neutral sentiment tweets, followed by positive sentiment tweets, and finally negative sentiment tweets.

2.23 Submission

Congrats, you just finished Project 1!

2.24 Submission

You're done!

Before submitting this assignment, ensure to:

1. Restart the Kernel (in the menubar, select Kernel->Restart & Run All)

2. Validate the notebook by clicking the "Validate" button

Finally, make sure to **submit** the assignment via the Assignments tab in Datahub