CSE199 Activity: Textual Sentiment Analysis Fall 2026

This week's activity is a 'get-acquainted' experience with: using Python, data-mining the web, sentiment analysis, and numerical interpretation of results. No past programming experience in Python is needed—the program is given and runs within (the free level of) a web app called Pickcode. Besides having fun, the goal is to raise discussion of data policy issues that supplement aspects of the third week's mini-project. This activity needs **Chrome** or **Firefox** or a downloaded Python system; the two browsers seem to work on any platform.

1 Introduction

Sentiment analysis means evaluating the emotions and possible intents of the authors of textual material. Here are two mainstream applications:

- Filtering communications and webpages for terrorist indications and gauges of threat levels.
- Evaluating customer (dis-)satisfaction with your company's products—and possibly diagnosing faults.

The presence of emotion distinguishes this from other textual modeling, such as the Signature *stylometry* system by which Peter Millican and co-workers identified the pseudonym Robert Galbraith as belonging to J.K. Rowling. The particular emotions are five personified in the 2015 Pixar movie *Inside Out*: *joy*, *anger*, *fear*, *disgust*, and *sadness*. [1]

We will use this tool to make simple measures of the "emotional temperature" of various webpages. One component of our measure is scientific. It is a human-scored (crowdsourced) *rating* of a lexicon of words for *intensity* in one or more of the five categories. For instance, the word "surprise" scores **0.172** for *fear* and **0.562** for *joy*. It scores zero for the other three categories. The plural "surprises" scores only for *joy*: **0.606**. The intensity is always less than 1; the most intense word for *anger* is "outraged": **0.964**. There are numerous NSFW terms in the lexicon—well, this activity is structured so that you won't see them unless you load a webpage that has them.

The lexicon is copyrighted by Dr. Saif Mohammed and the Canadian National Research Council (NRC). We are using an older version of his lexicon with permission from him and his co-worker Pierre Charron. Although the current version is freely downloadable from the project page, it carries terms of use (bottom of page) that forbid *re-distributing* it and were in effect for our version as well.

2 To Do Before Recitation

For background, please *skim-read* Dr. Mohammed's 10-page paper "Word Affect Intensities" (2018). You need not understand all the technical bits but should appreciate the following points:

- Even at professional level this kind of work is (and in 2026, still is) in early stages.
- This was almost the first lexicon with numerical rather than binary (i.e. 0-1 or yes/no) data on words for general purposes.

[1] The 2024 sequel Inside Out 2 features five new emotions: *anxiety*, *envy*, *embarrassment*, *ennui*, and *nostalgia*. Meanwhile, the creators of our tool and its training lexicon went in a different direction with categories of *anticipation*, *trust*, and *surprise*. We will stay with the basic five for simplicity.

- The numbers were computed (to three decimal places) not by whim but by some kind of regular scientific procedure based on examining large amounts of data and using "cloud and crowd" methods.
- The crowd-sourced numbers were checked against human annotators. They were found to correlate significantly highly with an average of human opinions. The correlation measures mentioned on page 2 are not the same as the "R^2" from the linear-model examples in lecture but are related.

Then finally please skim-read the Python code files `heatindex.py` and `heatlib.py`, which can be found in my https://www.cse.buffalo.edu/~regan/cse199/ folder. Note especially the lines with `location =` in `heatindex.py` and with `pageStr` and uses of `re.sub` in the `heatScore` function in `heatlib.py`, where there are alternatives to comment in or out. Those alternatives were needed with a web app called Python 3 Trinket but it has been discontinued as of 8/31/26. You are optionally welcome to run your code files in Google Colab or in your own Python 3 environment; the directions may vary as described below for Colab.

In brief, what the code does is tally all the lexicon words that appear in a loaded webpage. The score total of words in the *joy* category is multiplied by -1. The other categories each have multipliers of $+1$, which means their positive scores are just summed together. Adding the scores gives a *negative* result if there are more—and stronger—*joy* words than others; else a positive score. It may seem strange to you that joy is being modeled by a negative number, but ask yourself, to what emotions does *intensity* more often apply?

3 Setup Directions (15 minutes)

The runs are relatively simple and will be demo'ed in the week 2 Wednesday lecture. For Pickcode:

1. *Create a free account* at Pickcode; Click "View pricing," then under "Start with a free account" click "Sign up," say you're a student, give your age, and finally give e-mail and create a password.
2. You will see an empty "Sandbox." Click "New Project" and give a title ("HeatIndex" is fine). Select the language as Python, then under Project Type select "Console," and finally click "Create." You will see a split screen with black background.
3. Click the three bars at upper left next to `main.py`, then click "Add/Upload" at the bottom, and select "New blank file."
4. You will see "New File" at the top with three vertical dots to its right. Click those and select "Rename." It will let you type over the name: type `heatlib.py` all-lowercase. It is important to use that exact name. Finally click on the name and you will see two tabs above it, `heatlib.py` and `main.py` with the former highlighted.
5. Then go to Dr. Regan's CSE199 web folder to get the file `heatlib.py` — you can use https://cse.buffalo.edu/r̃egan/cse199/heatlib.py for a direct link. You should see the text of the file in your browser window. (If the file automatically downloads instead, see alternate directions below.)
6. Highlight all text (Ctrl-A or Cmd-A) and copy it (Ctrl-C or Cmd-C). Go back to the Pickcode window, and making sure the tab for `heatlib.py` is active, paste at line 1. (It is important that what you paste includes the Python open-comment line `'''` at the very top.)

7. Then click the tab for `main.py`, get the text of `heatindex.py` ffrom the web folder, and paste it *overwriting* the "Hello world!" line at the top. Do not change the name `main.py` and no file upload.

8. Finally click the standard green-triangle "play" icon at lower right. After a couple seconds, you should see a note about "disgust" being added as a category and the prompt `Paste a URL:` with a text box. Go to a website, copy and paste the URL, hit return, and you will get a lot of text and number results that you can scroll back through.

The old Trinket app would freeze if left untouched for some length of time. *Something similar happens with Pickcode.* You may get another input box with no screen feedback, or it may say "Starting code..." but then hang. Clicking the play button again seems *not* to work as it did with Trinket. The ultimate fix is to *refresh the entire webpage.*

If you type in a URL without the leading `https://` you may get a warning (on Trinket this would often fail). Some webpages may cause errors of three kinds:

- The webpage refuses the "GET" request in a way that causes an error. Restart and try another webpage.

- The webpage reports a decoding error. This will happen if you try to read a PDF or Word file or anything not in UTF-8 plain-HTML format. Restart and try another webpage.

- The webpage gives no output, or nondescript output. Twitter/X pages used to return only the string `x.com`. They now (2026) seem to return a scan of the user's most recent few posts, even without your being logged in to X. The *Wall Street Journal* used to give only a boiler-plate site index dump but now (2026) seems to include some of its top story. In previous years, some pages returned the text of an error saying "Something went wrong. Don't fret—let's give it another shot" and elsewhere, "Terms of service" and "Access control disabled." This matched the NRC lexicon on words including "wrong," "fret," "shot," "service," and "disabled," which at first seemed thematic for a news story covering a war. If you suspect this, comment-in the line `print(Ï read ;pageStr)` (on or around around line 180 of `heatlib.py`) to get a dump of the entire text received. Or just (restart and) try another webpage.

Basically the catchall for troubleshooting is "just restart and try another webpage."

Alternate directions with file upload (needed for Google Colab and optional with Pickcode): Download the files `heatlib.py` and `heatindex.py` to your machine with those names. (On Windows, right-click and select `Save link as...`) After getting a free account and logging in to Colab, click "New Notebook." You will see a mostly-blank screen with "Untitled1.ipynb" at the top—leave that name alone and click the folder icon at far left, second icon from the bottom. A sidebar "Files" will open up—wait a second and the standard upload icon with an up-arrow will appear below it. Use it to upload the two files one-by-one, ignoring any warning about their not being preserved.

In Colab, you then need to go to the "Start coding, or..." slot, type `%run heatindex.py` and hit the play icon. (Colab seems not to freeze; if you use it and encounter issues, please let us know.)

If you upload by name in Pickcode, you will see one of the names at lower right next to the play icon. If it is not already `heatindex.py`, click the name and you will get an option to select it. Then click the play icon as before.

4 Activity Directions

Get into your groups. The first part is to be done **individually** by everyone—for practice—but discussed briefly in groups. It also furthers the discussion about caveats to look for when loading webpages. The activity has a practice stage, then a game-playing stage, and a final "debrief."

Practice (10 minutes)

1. First, try my webpage https://cse.buffalo.edu/~regan/ Note that just a few words match the lexicon and the net score is only slightly positive. The words "rating inflation" are a chess term, but their presence could well make readers think of the economy, so their "affect" in dimensions of anger, fear, and sadness still applies.
2. Then try the webpage https://www.kimberlymarshall.com/ of Kimberly Marshall. The score should be in the vicinity of **-10**. Look at the words in the categories. It is interesting that "music" counts for both *joy* and *sadness*, and "musical" counts further for *anger*. But they score much higher for *joy*, so they augment the overall preponderance of "joy" terms.

Game-Playing Stage (10–15 minutes, including time to compare groups' results)

Your objective is to find *webpages with negative scores*—besides the one above. The group with the most-negative score wins brownie points. A second point can go to the group that finds the most webpages with negative scores—or the most diversity in such webpages. Think of kinds of webpages that would have incentive to generate an uplifted impression on their readers.

The following materials are to be logged and submitted by each group:

1. URLs and brief descriptions of webpages you try.
2. For at least two pages with negative scores, a discussion of why the page would be interested in promoting *joy*.
3. A brief reaction to the following "skeptical debrief"—which can be put into your individual final feedback rather than a group submission (and so won't be graded as part of the activity itself).

The grading rubric is: **1 pt.** for participation, **1 pt.** for showing webpages you tried, and **1 pt.** for finding negative scores and discussing why. Especially the final one-point can be divided as **0.25** for each page with a negative score and **0.25** for the accompanying discussion.

5 Skeptical Debrief (as time allows, or offline)

I hope the webpage scoring is fun, maybe even addictive... But before we get carried away, we need a "coach's huddle" to address the general question of **whether this is responsible data science**.

The most important thing to know about the numerical results given at the very end is that they are **not designed by Dr. Mohammad** but rather by me, KWR. This project design performs a "reduction" on his work. The final number comes from adding up the 'anger,' 'fear,' 'disgust', and 'sadness' scores and *subtracting* the total 'joy' score. Then my code divides by the total number of words read and—totally arbitrarily for better "eye candy"—multiplies that by 1,000. In spot-tests,

this puts the final numbers roughly on a familiar 0-to-10 scale, concentrated more near zero, and possibly negative—when 'joy' outweighs the other four categories. So it has some street sense. But it's still doing "push a button and get a number"—which is what everybody wants in order to make their jobs simpler but which comes with blinkers and dangers.

Put another way, I have slapped a layer of "multipliers" onto Dr. Mohammad's model. The default multipliers pass the "Occam's Razor" test of being simple: +1 for each *anger*, *disgust*, *fear*, and *sadness* word, and −1 for *joy*. The code allows you to change them at will, but they are still *arbitrary*—that is, not justified by theoretical considerations and/or empirical testing.

If the debrief gets to this point in the recitation session, that's fine. If there is further time, then we can go into the following—or if not, invite these further considerations offline. They discuss *what* we want to model, and whether the "multipliers" bolted onto Dr. Mohammed's lexicon and the concept of a single "heat score" make some undue presumptions:

- They presume we are trying to model 'light'-versus-'dark' in some way.
- If we want to model *intensity* of any kind, we should use +1 for *joy* too.
- At least the +1 multipliers are leaving Dr. Mohammad's carefully-crafted numbers alone...
- But adding them up and dividing by the total number of words is still "reducing" his work. What is that ratio supposed to represent?
- If we append to the bottom of a page an equal amount of completely neutral text, we will *halve* the score. Does this make sense—shouldn't we give more weight to text at the top?
- Perhaps 'anger' should be regarded as generating more "heat" than 'fear' (which is more passive) and certainly 'sadness.' (Indeed, those who have seen *Inside Out* may recall its upshot that the "negative" emotions have positive roles.)
- Try multipliers of +10 for anger, +6 for fear, +2 for sadness and (really "winging it" now) −5 for joy (disgust may be left at +1). Repeat some earlier URLs. Do any scores shift notably? (Note: Just doubling the multipliers will not double the score, because the code has a further "normalizing" division by the sum of the absolute values of the multipliers.)
- Try other multipliers too. Each combo constitutes a different *model* of "webpage intensity." Then reflect on the most important question:

> From the myriad parameter combinations, how can we determine **one** that is *correct*, *best*, or at least *neutrally justifiable*? In particular, how could we responsibly *train* the model?

Answering this question first requires determining what we want to model and what the purpose is. Suppose we adopt the "light-versus-dark" purpose. The first thing we might try is to identify a corpus of pages that represent the "neutral" middle. Then we want to define our scale and train our multipliers so that those pages get a score of 0. This still does only a quarter of the work we need—in technical jargon, it leaves three "degrees of freedom" in the four parameters we are fitting.

Further progress needs asking, what are we trying to *predict*? Here we might come full circle to the "threat levels" application mentioned at the outset. Suppose we have a corpus of pages that were reliably associated with the same threat level in past history. Some pages might have more 'anger,' others more 'fear' and so on. We can train our multipliers to equalize those pages.

A general methodology emerges from the idea that how morose or bubbly a webpage written today is can predict how morose or bubbly a page written by the same person(s) tomorrow will be. Or we can apply this prediction idea from sentence to sentence or phrase to phrase. Now we are in the world of those N-grams. Predicting a probability range on the nature of the $(N+1)$st item from the previous N in a continuing series involves building a so-called *Markov model* (or *Markov chain*, named for the Russian mathematician Andrey Andreyevich Markov). The Markov model M can use the principle that words continuing the score trend of the previous N are most likely. Its results thus depend on the scores and hence on the multipliers we choose. The likelihood principle in this context yields the training policy:

Find the combination of parameters that maximizes the probability that M projects for the record of what actually happened.

This sounds like "making yourself look good by predicting the past" but the power of this principle carries into the future. It still, however, takes quite a bit of work to build M and then some computer muscle to carry out the *maximum likelihood estimation* (MLE). In fact, the architecture of Large Language Models (LLMs) and their underlying **transformer** mechanism with **attention**, enshrines even more the principle:

Predict the past to generate the future.

If this is taken as the *only* directive, then it leaves out human-welfare qualities we desire for the future. It will at most pick up qualities from the past, which have often worked against human welfare.

But even before we consider responsibility for AI training, we have to confront basic responsibility issues in this sentiment application from the above discussion. The fact that the results are plausible even in this simplistic formulation does not confer a scientific imprimatur. It is if anything more prone to appeal to your incoming biases than to reflect the application of professional standards that we seek in the coming years to confer on you at UB.